arXiv:1610.05954v1 [cs.PL] 19 Oct 2016

Dissertation:

Unfolding Semantics of the Untyped λ-Calculus with letrec

Jan Rochel <jan@rochel.info>

defended on June 20, 2016
online version, revision: October 20, 2016

Dedication: *Gewidmet meiner ehemaligen Klavierlehrerin Eva-Maria Rieckert, die mir Musik auf eine Weise vermittelte, die mich bis heute über das Musizieren hinaus prägt.*

Promotoren:

Prof.dr. S. D. Swierstra
Prof.dr. V. van Oostrom

Copromotor:

Dr. C. Grabmayer

Unfolding Semantics of the Untyped λ-Calculus with letrec

Ontvouwingssemantiek van de ongetypeerde λ-calculus met letrec
(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht
op gezag van de rector magnificus, prof.dr. G.J. van der Zwaan,
ingevolge het besluit van het college voor promoties in het openbaar
te verdedigen op maandag 20 juni 2016 des middags te 2.30 uur

door

Jan Rochel

geboren op 20 juli 1984
te Bretten, Duitsland

Dit proefschrift werd (mede) mogelijk gemaakt met financiële steun van de Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO).

Contents

Contents **iii**

Preface **v**
Abstract vi

0 $\lambda_{\mathsf{letrec}}$ and Unfolding **1**
0.1 Introduction 2
0.2 λ-terms and $\lambda_{\mathsf{letrec}}$-terms – informal 5
0.3 Unfolding $\lambda_{\mathsf{letrec}}$-terms – informal 6
0.4 Preliminaries 11
0.5 λ-terms and $\lambda_{\mathsf{letrec}}$-terms – CRS formalisation 12
0.6 Unfolding Rules – CRS formalisation 13
0.7 Unfolding Semantics of $\lambda_{\mathsf{letrec}}$ 18

1 Expressibility in $\lambda_{\mathsf{letrec}}$ **23**
1.1 Overview 23
1.2 Introduction 24
1.3 Preliminaries 31
1.4 Regular and strongly regular λ-terms 38
1.5 Observing $\lambda_{\mathsf{letrec}}$-terms by their generated subterms 58
1.6 Proving regularity and strong regularity 65
1.7 Binding–Capturing Chains 89
1.8 Expressibility by terms in $\boldsymbol{\lambda}_{\mathsf{letrec}}$ 99
1.9 λ-transition-graphs 108
1.10 Summary 112

2 Term Graph Representations for Strongly Regular λ-Terms **114**
2.1 Overview 114

2.2 Preliminaries . . . 115
2.3 Introduction . . . 120
2.4 λ-higher-order-Term-Graphs . . . 122
2.5 Abstraction-prefix based λ-ho-term-graphs . . . 127
2.6 λ-Term-Graphs without Scope Delimitiers . . . 131
2.7 λ-Term-Graphs with Scope Delimiters . . . 134
2.8 Not closed under (functional) bisimulation . . . 144
2.9 Closed under functional bisimulation . . . 146
2.10 Transfer of the complete-lattice property to λ-ho-term-graphs . . . 156
2.11 Summary . . . 160

3 Maximal Sharing in $\boldsymbol{\lambda}_{\mathsf{letrec}}$ 162
3.1 Overview . . . 162
3.2 Preliminaries . . . 163
3.3 Introduction . . . 163
3.4 Overview: Methods and Formalisms . . . 165
3.5 Interpretion of $\lambda_{\mathsf{letrec}}$-terms as λ-ap-ho-term-graphs . . . 168
3.6 Interpretion of $\lambda_{\mathsf{letrec}}$-terms as λ-term-graphs . . . 178
3.7 Readback of λ-term-graphs . . . 182
3.8 Complexity analysis . . . 192
3.9 Implementation . . . 194
3.10 Modifications, extensions and applications . . . 194

Bibliography 199

A Examples: Graph Translation 205

B Implementation Showcase 215

C Confluence of unfolding $\lambda_{\mathsf{letrec}}$-terms 216

D Unfolding with a Single Rule 232

Curriculum Vitae 237

Samenvatting in het Nederlands (Summary in Dutch) 238

Lay Summary 240

Acknowledgements 243

Preface

This thesis documents research which was carried out as part of an NWO[1] research project under my promotors Vincent van Oostrom and Doaitse Swierstra titled 'Realising Optimal Sharing'. The objective of the project was to investigate whether the theory of 'optimal evaluation of the λ-calculus' could be used in practise to increase the execution efficiency for programs written in functional programming languages. As regards the original project goal my research was unsuccessful. This is not due to a lack of results but due to a distraction: when approaching the subject of 'optimal sharing' – which is 'dynamic' in the sense that it concerns sharing that an evaluator maintains at run time – more fundamental, still unresolved questions emerged concerning 'static' sharing – which is the sharing inherent in a given program definition. Therefore the presented results are not about optimal and therefore 'dynamic' sharing but solely about 'static' sharing.

The research was carried out in close collaboration with Clemens Grabmayer. All of the fundamental ideas and results are due to joint efforts and fruitful – if sometimes fierce – debate. We complemented each other splendidly: I profited much from Clemens Grabmayer's expertise with formal systems, while I myself could contribute my proficiency with functional programming languages and compiler construction; also I have worked out an implementation of our methods. The substance of this thesis is from three research papers [19, 24, 22] published in the context of my doctoral research with Grabmayer and Rochel as authors, presented here in a more coherent narrative with supplementary explanations and examples. In order to do justice to the Clemens Grabmayer's substantial contribution, authors who wish to cite this thesis in their work are kindly advised to cite at least one of the papers alongside.

This document is structured according to these three papers: chapter 0

[1] Nederlandse Organisatie voor Wetenschappelijk Onderzoek

introduces the $\lambda_{\mathsf{letrec}}$-formalism and our rewriting system for unfolding $\lambda_{\mathsf{letrec}}$-terms. There we also pose the problems that are resolved in the following three chapters, chapter 1, chapter 2, and chapter 3, each of which corresponds to one of these papers. Note, that the formalisms in this thesis deviate to varying degrees from their original form in the papers. The changes were required for the formalims to be consistent throughout the chapters.

Abstract

In this thesis we investigate the relationship between finite terms in $\boldsymbol{\lambda}_{\mathsf{letrec}}$, the λ-calculus with letrec, and the infinite λ-terms they express. We say that a $\lambda_{\mathsf{letrec}}$-term *expresses* a λ-term if the latter can be obtained as an infinite unfolding of the former. Unfolding is the process of substituting occurrences of function variables by the right-hand side of their definition. We consider the following questions:

(i) How can we characterise those infinite λ-terms that are $\lambda_{\mathsf{letrec}}$-expressible?

(ii) given two $\lambda_{\mathsf{letrec}}$-terms, how can we determine whether they have the same unfolding?

(iii) given a $\lambda_{\mathsf{letrec}}$-term, can we find a more compact version of the term with the same unfolding?

To tackle these questions we introduce and study the following formalisms:

- a rewriting system for unfolding $\lambda_{\mathsf{letrec}}$-terms into λ-terms
- a rewriting system for 'observing' λ-terms by dissecting their term structure
- higher-order and first-order graph formalisms together with translations between them as well as translations from and to $\boldsymbol{\lambda}_{\mathsf{letrec}}$

We identify a first-order term graph formalism on which bisimulation preserves and reflects the unfolding semantics of $\boldsymbol{\lambda}_{\mathsf{letrec}}$ and which is closed under functional bisimulation. From this we derive efficient methods to determine whether two terms are equivalent under infinite unfolding and to compute the maximally shared form of a given $\lambda_{\mathsf{letrec}}$-term.

Chapter 0

$\lambda_{\mathsf{letrec}}$ and Unfolding

§ 0.0.1 (abstract)**.** This thesis concerns itself with terms in the λ-calculus with letrec, specifically with *unfolding* these terms. Unfolding refers to the process of substituting occurrences of let-bound variables by their definition. We define unfolding by means of a rewriting system. We study the properties of that rewriting system and build various formal systems on top of it to derive further results. These results include:

chapter 1: a characterisation of the infinite λ-terms that can be expressed finitely as $\lambda_{\mathsf{letrec}}$-terms.

chapter 2: a graph representation for $\lambda_{\mathsf{letrec}}$-expressible λ-terms.

chapter 3: practical and efficient methods for transforming a $\lambda_{\mathsf{letrec}}$-term into a maximally compact form; and deciding whether two $\lambda_{\mathsf{letrec}}$-terms have the same unfolding.

§ 0.0.2 (required background)**.** The reader is expected to have some proficiency with functional programming languages based on the λ-calculus [12, 6] (preferably Haskell), which is most likely required to understand the presented results and their relevance. Furthermore, the formal systems used for reasoning in this thesis are mostly rewriting systems. Therefore, at least a basic background in term rewriting [51] is assumed.

§ 0.0.3 (chapter overview)**.** This chapter gives an overview over $\lambda_{\mathsf{letrec}}$ and outlines the perspective from which we will study this calculus. We provide definitions and basic properties of the rewriting system by which we unfold

$\lambda_{\mathsf{letrec}}$-terms. Then we will pose the questions and problems to contextualise the results presented in the following chapters.

0.1 Introduction

§ 0.1.1 ($\lambda_{\mathsf{letrec}}$ as an abstraction of functional programming languages)**.** The λ-calculus is a formal system in computer science and logic for expressing computation. It is the model of computation at the core of functional programming languages. In this thesis we will look at one specific instance of the λ-calculus, namely the untyped λ-calculus extended by the letrec-construct, or in short $\lambda_{\mathsf{letrec}}$. In that form, it serves well as a minimalistic abstraction of functional programming languages like Haskell. While Haskell is a typed language, it is typically translated into a simplified form during the compilation process in which type information is discarded (*type erasure*). Thus types can be regarded as auxiliary means for the programmer, and can be neglected when looking at the evaluation semantics of a type-checked program.

§ 0.1.2 (the letrec-construct)**.** The letrec-construct serves a number of purposes. By allowing us to bind subterms to variables it adds to the simple, untyped λ-calculus the means for:

modularisation: Subterms with a specific purpose can be given a descriptive name and more easily be treated as entities of their own.

sharing: Instead of repeating identical subterms at different locations, a subterm can be defined once and be referenced by its function definition multiple times.

cyclicity: Function definitions can contain references to themselves, allowing for cyclic (and mutually cyclic) bindings.

While all the above can also be achieved by a collection of top-level bindings, the letrec has one distinguishing characteristic, which is that function bindings can be defined at any position in the term. The scope of thus defined function bindings does not range over the entire program but can only be used underneath the position of their definition. This locality of definitions allows for a more structured approach to programming. Also, it can refer to λ-variables bound outside of the binding which results in fewer β-reduction steps during evaluation.

Before we concern ourselves with formal definitions let us first fix some notation and terminology.

Notation 0.1.3 (let = letrec)**.** Throughout this thesis we write let to denote the letrec-construct, as is done in Haskell.

Terminology 0.1.4 (let-expression)**.** A term that starts with a let, thus a term that has the form let B in L, is called a let-*expression*.

Terminology 0.1.5 (binding group)**.** We call the collection of bindings B defined in a let-expression let B in L a *binding group*.

Terminology 0.1.6 (function binding, function variable, let-bound variable)**.** We call the equations of a let-expression *function bindings* or simply *bindings*. We call the variable on the left-hand side of a binding a let-*bound variable*, or a *function variable*.[1]

Terminology 0.1.7 (body)**.** The part L of the expression let B in L is called the *body* of the let-expression.

Example 0.1.8. Consider the $\lambda_{\mathsf{letrec}}$-term let fix = $\lambda f.\, f\ (\mathrm{fix}\ f)$ in fix. The binding group of the let-expression consists of a single function binding fix = $\lambda f.\, f\ (\mathrm{fix}\ f)$ which binds the term $\lambda f.\, f\ (\mathrm{fix}\ f)$ to the function variable fix. The body of the let-expression consists of an occurrence of the function variable fix.

Remark 0.1.9 (Turing completeness, well-typedness, termination)**.** Typed λ-calculi (with finite types) are strongly normalising, which means that every computation in such a calculus terminates. Therefore such calculi are not Turing complete. However, strong normalisation does not hold for typed λ-calculi that allow for cyclic definitions, such as $\lambda_{\mathsf{letrec}}$. Therefore the letrec can also be seen as a way to restore Turing completeness for a typed λ-calculus. Other approaches would be fixed-point combinators or top-level bindings. The letrec offers the most convenience from a programmer's point of view.

§ 0.1.10 (infinite unfolding)**.** A $\lambda_{\mathsf{letrec}}$-term L can be seen as a finite representation of a (possibly) infinite λ-term M, which we obtain by repeatedly substituting every occurrence of a function variable by the right-hand side of the corresponding binding. We call M the *infinite unfolding* of L, and we write $[\![L]\!]_{\boldsymbol{\lambda}} = M$. A definition for $[\![\cdot]\!]_{\boldsymbol{\lambda}}$ is given later (definition 0.7.1).

[1] While a term bound by let to a variable may well be constant (i.e. not a λ-abstraction) we still call such a binding a function binding and the let-bound variable a function variable.

Example 0.1.11 (infinite unfolding of fix id)**.** Let us consider a naive implementation of the fix-function applied to $\lambda x.\, x$, the identity function:

$$\begin{aligned}&[\![(\mathsf{let}\ \mathrm{fix} = \lambda f.\, f\ (\mathrm{fix}\ f)\ \mathsf{in}\ \mathrm{fix})\ (\lambda x.\, x)]\!]_{\boldsymbol{\lambda}}\\ &\quad = (\lambda f.\, f\ ((\lambda f.\, f\ (\ldots\ f))\ f))\ (\lambda x.\, x)\end{aligned}$$

Notation 0.1.12 (ellipsis: $\ldots$)**.** Note that (as in the example above) we will be using the ellipsis as an informal notation for 'and so on' extensively. It occurs both in infinite terms and infinite rewriting sequences. Instead of providing a first-order formalisation of the ellipsis, we trust that the reader will find the context always sufficient to infer the shape of the entire term or rewriting sequence.

§ 0.1.13 (evaluation and unfolding)**.** While semantics of $\lambda_{\mathsf{letrec}}$-based programming languages can be defined via infinite unfolding, evaluation of programs on a computer cannot operate on infinite terms but must rely on a finite representation. The evaluation of such programs requires in addition to β-reduction a mechanism to unfold let-expressions. This is best modelled in a rewriting system that extends the λ-calculus by unfolding rules.

Example 0.1.14 (leftmost-outermost evaluation of fix id)**.** Let us evaluate a small example program to see how unfolding comes into play in the course of evaluating $\lambda_{\mathsf{letrec}}$-terms:

$$(\mathsf{let}\ \mathrm{fix} = \lambda f.\, f\ (\mathrm{fix}\ f)\ \mathsf{in}\ \mathrm{fix})\ (\lambda x.\, x)$$

This term has no (visible) β-redex as the $\lambda f. \ldots$-abstraction is 'blocked' by the surrounding let. In order to turn it into a proper β-redex we need to unfold its definition, which means essentially substituting the right-hand side of the binding for both occurrences of the function variable fix:

$$\twoheadrightarrow_{\triangledown} \quad (\lambda f.\, f\ (\mathsf{let}\ \mathrm{fix} = \lambda f.\, f\ (\mathrm{fix}\ f)\ \mathsf{in}\ \mathrm{fix})\ f)\ (\lambda x.\, x)$$

In the formalisation of unfolding as a rewriting system (see section 0.3) this requires actually a number of steps, hence the many-step reduction $\twoheadrightarrow_{\triangledown}$. Now that we have a visible β-redex to contract, we can substitute $\lambda x.\, x$ for f:

$$\rightarrow_{\beta} \quad (\lambda x.\, x)\ ((\mathsf{let}\ \mathrm{fix} = \lambda f.\, f\ (\mathrm{fix}\ f)\ \mathsf{in}\ \mathrm{fix})\ (\lambda x.\, x))$$

Next, we apply the identity function and we arrive at the initial term and evaluation continues as above and will thus never terminate.

$$\rightarrow_{\beta} \quad (\mathsf{let}\ \mathrm{fix} = \lambda f.\, f\ (\mathrm{fix}\ f)\ \mathsf{in}\ \mathrm{fix})\ (\lambda x.\, x)\ \twoheadrightarrow_{\triangledown}\ \ldots$$

Remark 0.1.15 ($\triangledown$)**.** Using the triangle as a symbol for unfolding is inspired by graph rewriting systems like Lambdascope [47] where sharing is indicated by an explicit sharing node in the shape of a triangle, with multiple incoming edges at the top (shared occurrences) and one outgoing edge at the bottom (to the shared subgraph). An unsharing step would then 'unzip' the shared subgraph node by node, with the triangle acting as the zipper foot.

§ 0.1.16 (mixing unfolding and β-reduction)**.** A simple interpreter for $\lambda_{\mathsf{letrec}}$ proceeds along the lines of example 0.1.14, i.e. by interspersing β-reduction with unfolding steps in a combined rewriting system. Most of the scientific works around unfolding $\lambda_{\mathsf{letrec}}$-terms are works that study evalutors that include unfolding rules with β-reduction (and possibly α-reduction) in a rewriting system.

§ 0.1.17 (unfolding as a subject of study)**.** This thesis however focusses entirely on the unfolding portion of the semantics of $\lambda_{\mathsf{letrec}}$; β-reduction will henceforth play a marginal role at most.[2]

§ 0.1.18 (outlook)**.** In the following two sections we will define the term language for $\lambda_{\mathsf{letrec}}$-terms as well as a rewriting system for unfolding terms in $\lambda_{\mathsf{letrec}}$. First we give an informal account to introduce the notation that we will be using for examples. Afterwards we provide sound formalisations.

0.2 λ-terms and $\lambda_{\mathsf{letrec}}$-terms – informal

§ 0.2.1 (overview)**.** This section provides first-order notations for terms in the λ-calculus and the $\lambda_{\mathsf{letrec}}$-calculus. Mind, that these are not formal definitions and are only used to convey an intuition for the issue at hand. Only in section 0.5 the definitions are formalised in the CRS (Combinatory Reduction System) framework.

§ 0.2.2 (set of λ-terms)**.** Let V be a set of variable names. The set of λ-terms is 'coinductively' defined by the following grammar, where $x \in V$:

$$\begin{array}{llll} \text{(term)} & L & ::= & \lambda x.\, L \quad \text{(abstraction)} \\ & & \mid & L\ L \quad \text{(application)} \\ & & \mid & x \quad \text{(variable)} \end{array}$$

[2] This vaguely suggests possible future research: it could be a promising venture to develop modular semantics for $\boldsymbol{\lambda}_{\mathsf{letrec}}$, i.e. '$\beta$-reduction modulo unfolding', and express existing evaluators in terms of this semantis in a modular manner.

§ 0.2.3 (infinite terms)**.** Mind, that we interpret this grammar coinductively (i.e. as a final coalgebra). Therefore finite as well as infinite terms arise from it. Since unfoldings of $\lambda_{\mathsf{letrec}}$-terms are typically infinite, in this thesis we will more often than not deal with infinite λ-terms. $\lambda_{\mathsf{letrec}}$-terms on the other hand will always be finite.

Example 0.2.4 (a simple infinite λ-term)**.** See example 0.1.11.

Adding a production for the letrec-construct to the above grammar, we obtain a grammar for $\lambda_{\mathsf{letrec}}$.

§ 0.2.5 (set of $\lambda_{\mathsf{letrec}}$-terms)**.** Let V be a set of variable names. The set of $\lambda_{\mathsf{letrec}}$-terms is inductively defined by the following grammar, where $x, f_1, \ldots, f_n \in V$:

(term)	L	::=	$\lambda x.\, L$	(abstraction)
		|	$L\ L$	(application)
		|	x	(variable)
		|	$\mathsf{let}\ B\ \mathsf{in}\ L$	(letrec)
(binding group)	B	::=	$f_1 = L, \ldots, f_n = L$	(bindings)
			($f_1, \ldots, f_n$ all distinct)	

0.3 Unfolding $\lambda_{\mathsf{letrec}}$-terms – informal

On this grammar we will now develop a rewriting system to describe unfolding of $\lambda_{\mathsf{letrec}}$-terms in an informal notation.

§ 0.3.1 (substitution of function variables)**.** First we need a rule to perform the actual unfolding, i.e. the substitution of a function variable occurrence by the right-hand side of its definition.

$$\mathsf{let}\ B_1,\ f = L,\ B_2\ \mathsf{in}\ f \quad \to_{\mathrm{rec}} \quad \mathsf{let}\ B_1,\ f = L,\ B_2\ \mathsf{in}\ L$$

This rule is only applicable to a function binding with a function variable as its body.

§ 0.3.2 (distributing function bindings)**.** In case of a more complex body, we distribute the function binding over its constituents. The following two rules distribute function bindings over applications and abstractions.

$$\begin{aligned} \mathsf{let}\ B\ \mathsf{in}\ L_0\ L_1 &\quad \to_{@} \quad (\mathsf{let}\ B\ \mathsf{in}\ L_0)\ (\mathsf{let}\ B\ \mathsf{in}\ L_1) \\ \mathsf{let}\ B\ \mathsf{in}\ \lambda x.\, L &\quad \to_{\lambda} \quad \lambda x.\, \mathsf{let}\ B\ \mathsf{in}\ L \end{aligned}$$

§ 0.3.3 (merging function bindings)**.** Two nested function bindings can be merged into one:

$$\mathsf{let}\ B_0\ \mathsf{in}\ \mathsf{let}\ B_1\ \mathsf{in}\ L \quad \rightarrow_{\mathrm{letrec}} \quad \mathsf{let}\ B_0, B_1\ \mathsf{in}\ L$$

§ 0.3.4 (name clashes, α-renaming)**.** Note, that the rules above are all in informal notation. The actual definitions (section 0.6) are CRS rewriting rules. Thus, name clashes (as for instance two functions of the same name being defined in B_0 as well as in B_1 in the above rewriting rule) are not a problem that we need to concern ourselves with, as they are dealt with by the CRS formalism. When using first-order notation we will rename variables whenever necessary (or convenient).

§ 0.3.5 (garbage collection)**.** The above rules would suffice to distribute the function bindings to the corresponding function variable occurrences and unfold them. In order to obtain an unfolded λ-term without any residual function bindings, we include these garbage-collection rules:

$$\mathsf{let}\ f_1 = L_1 \ldots f_n = L_n\ \mathsf{in}\ L \quad \rightarrow_{\mathrm{red}} \quad \mathsf{let}\ f_{j_1} = L_{j_1} \ldots f_{j_{n'}} = L_{j_{n'}}\ \mathsf{in}\ L$$
(if $f_{j_1}, \ldots, f_{j_{n'}}$ are the function variables reachable from L)

$$\mathsf{let}\ \mathsf{in}\ L \quad \rightarrow_{\mathrm{nil}} \quad L$$

The latter discards empty function bindings, while the former removes all function bindings from a function binding that are not 'reachable'. We consider a function binding to be 'reachable' if the corresponding function variable either occurs in the body of the let-expression or in any other of the function bindings that is 'reachable'. The side condition, which ensures that only superfluous bindings are removed from the binding group is non-trivial and requires a reachability analysis because there might be mutually recursive unused function bindings.

The above rules define a rewriting system for unfolding $\lambda_{\mathsf{letrec}}$-terms to (possibly infinite) λ-terms.

Example 0.3.6 (unfolding derivation of fix = $\lambda f.\,\mathsf{let}\ r = f\ r\ \mathsf{in}\ r$)**.**

$$\begin{aligned}
& \lambda f.\,\mathsf{let}\ r = f\ r\ \mathsf{in}\ r \\
\to_{\mathrm{rec}}\ & \lambda f.\,\mathsf{let}\ r = f\ r\ \mathsf{in}\ f\ r \\
\to_{@}\ & \lambda f.\,(\mathsf{let}\ r = f\ r\ \mathsf{in}\ f)\ (\mathsf{let}\ r = f\ r\ \mathsf{in}\ r) \\
\to_{\mathrm{red}}\ & \lambda f.\,(\mathsf{let}\ \ \mathsf{in}\ f)\ (\mathsf{let}\ r = f\ r\ \mathsf{in}\ r) \\
\to_{\mathrm{nil}}\ & \lambda f.\,f\ (\mathsf{let}\ r = f\ r\ \mathsf{in}\ r)
\end{aligned}$$

and therefore:

$$\begin{aligned}
& \lambda f.\,\mathsf{let}\ r = f\ r\ \mathsf{in}\ r \\
\twoheadrightarrow_{\triangledown}\ & \lambda f.\,f\ (\mathsf{let}\ r = f\ r\ \mathsf{in}\ r) \\
\twoheadrightarrow_{\triangledown}\ & \lambda f.\,f\ (f\ (\mathsf{let}\ r = f\ r\ \mathsf{in}\ r)) \\
\twoheadrightarrow_{\triangledown}\ & \lambda f.\,f\ (f\ (f\ \ldots))
\end{aligned}$$

We say that fix *unfolds to* $\lambda f.\,f\ (f\ (f\ \ldots))$ and write

$$[\![\mathrm{fix}]\!]_{\boldsymbol{\lambda}} = \lambda f.\,f\ (f\ (f\ \ldots))$$

§ 0.3.7 (meaningless bindings)**.** However, not every $\lambda_{\mathsf{letrec}}$-term represents a λ-term. For instance the $\lambda_{\mathsf{letrec}}$-term L defined as

$$\lambda x.\,\mathsf{let}\ f = f\ \mathsf{in}\ f\ x$$

has a meaningless function binding $f = f$ that does not unfold to a λ-term. The rewriting rules above admit only the cyclic rewriting sequence $L \to_{\mathrm{rec}} L$. Therefore $L \notin \mathrm{dom}([\![\cdot]\!]_{\boldsymbol{\lambda}})$, which means that the unfolding semantics $[\![\cdot]\!]_{\boldsymbol{\lambda}}$ based on these rules can only be partial. In order to obtain a total unfolding semantics, we include a constant symbol to signified that the term is undefined at the point of its occurrence. As is customary since [2] we use the 'black hole' symbol $\bullet$, which we include in an extended version of the grammars for λ-terms and $\lambda_{\mathsf{letrec}}$-terms. The unfolding semantics of L will then be $\lambda x.\,\bullet\ x$. On the extended grammar we define two additional rules for turning meaningless bindings into black holes:

$$\begin{aligned}
&\mathsf{let}\ B_1, f = g, B_2\ \mathsf{in}\ L \quad \to_{\mathrm{tighten}} \quad \mathsf{let}\ B_1[f := g], B_2[f := g]\ \mathsf{in}\ L[f := g] \\
&\qquad (\text{if } g \text{ is defined in } B_1 \text{ or } B_2) \\
&\mathsf{let}\ B_1, f = f, B_2\ \mathsf{in}\ L \quad \to_{\bullet} \qquad \mathsf{let}\ B_1[f := \bullet], B_2[f := \bullet]\ \mathsf{in}\ L[f := \bullet]
\end{aligned}$$

The former rule inlines alias functions, which simplifies meaningless bindings to the form $f = f$, such that they can be turned into a black hole by the latter rule. Thus we obtain an extended unfolding semantics $[\![\cdot]\!]_{\boldsymbol{\lambda\bullet}}$ which (in contrast to $[\![\cdot]\!]_{\boldsymbol{\lambda}}$) is defined for all $\lambda_{\mathsf{letrec}}$-terms.

Example 0.3.8 ($\lambda_{\mathsf{letrec}}$-term with a meaningless binding)**.** The rules for handling meaningless bindings allow us to reduce L from above to a normal form, which means that $L \in \mathrm{dom}([\![\cdot]\!]_{\boldsymbol{\lambda\bullet}})$, or in particular $[\![\lambda x.\, \mathsf{let}\ f = f\ \mathsf{in}\ f\ x]\!]_{\boldsymbol{\lambda\bullet}} = \lambda x.\, \bullet\ x$ as is witnessed by the following rewriting sequence:

$$
\begin{array}{ll}
 & \lambda x.\, \mathsf{let}\ f = f\ \mathsf{in}\ f\ x \\
\to_{@} & \lambda x.\, (\mathsf{let}\ f = f\ \mathsf{in}\ f)\ (\mathsf{let}\ f = f\ \mathsf{in}\ x) \\
\to_{\mathrm{red}} & \lambda x.\, (\mathsf{let}\ f = f\ \mathsf{in}\ f)\ (\mathsf{let}\ \ \mathsf{in}\ x) \\
\to_{\mathrm{nil}} & \lambda x.\, (\mathsf{let}\ f = f\ \mathsf{in}\ f)\ x \\
\to_{\bullet} & \lambda x.\, \mathsf{let}\ \ \mathsf{in}\ \bullet\ x \\
\to_{\mathrm{nil}} & \lambda x.\, \bullet\ x
\end{array}
$$

Example 0.3.9 (meaninglessness due to mutual recursion)**.** $\mathsf{let}\ f = g,\ g = f\ \mathsf{in}\ f$ is meaningless due to mutually recursive functions. With the aid of $\to_{\mathrm{tighten}}$ the term can be reduced to a normal form:

$$
\begin{array}{ll}
 & \mathsf{let}\ f = g,\ g = f\ \mathsf{in}\ f \\
\to_{\mathrm{tighten}} & \mathsf{let}\ f = g,\ g = g\ \mathsf{in}\ g \\
\to_{\mathrm{red}} & \mathsf{let}\ g = g\ \mathsf{in}\ g \\
\to_{\bullet} & \mathsf{let}\ \ \mathsf{in}\ \bullet \\
\to_{\mathrm{nil}} & \bullet
\end{array}
$$

Example 0.3.10 (meaninglessness due to nested mutual recursion)**.** Let us consider another $\lambda_{\mathsf{letrec}}$-term L defined as $\mathsf{let}\ \ f = \mathsf{let}\ g = f\ \mathsf{in}\ g\ \ \mathsf{in}\ \ f$, which illustrates that meaninglessness is not always tied to a simple pattern. $L \notin \mathrm{dom}([\![\cdot]\!]_{\boldsymbol{\lambda}})$, but $L \in \mathrm{dom}([\![\cdot]\!]_{\boldsymbol{\lambda\bullet}})$ as is witnessed by the following rewriting

sequence:

$$\begin{array}{ll} & \mathsf{let}\ f = \mathsf{let}\ g = f\ \mathsf{in}\ g\ \mathsf{in}\ f \\ \to_{\mathrm{rec}} & \mathsf{let}\ f = \mathsf{let}\ g = f\ \mathsf{in}\ f\ \mathsf{in}\ f \\ \to_{\mathrm{red}} & \mathsf{let}\ f = \mathsf{let}\ \ \mathsf{in}\ f\ \mathsf{in}\ f \\ \to_{\mathrm{nil}} & \mathsf{let}\ f = f\ \mathsf{in}\ f \\ \to_{\bullet} & \mathsf{let}\ \ \mathsf{in}\ \bullet \\ \to_{\mathrm{nil}} & \bullet \end{array}$$

§ 0.3.11 ($\to_{\triangledown}$ and $\to_{\triangledown_\bullet}$)**.** We define rewriting relations $\to_{\triangledown}$ and $\to_{\triangledown_\bullet}$ for unfolding $\lambda_{\mathsf{letrec}}$-terms, where $\to_{\mathrm{tighten}}$ and $\to_{\bullet}$ are included only in $\to_{\triangledown_\bullet}$:

$$\begin{array}{lcl} \to_{\triangledown} & = & \bigcup\{\to_\rho \mid \rho \in \{@, \lambda, \mathrm{letrec}, \mathrm{red}, \mathrm{nil}\}\} \\ \to_{\triangledown_\bullet} & = & \bigcup\{\to_\rho \mid \rho \in \{@, \lambda, \mathrm{letrec}, \mathrm{red}, \mathrm{nil}, \mathrm{tighten}, \bullet\}\} \end{array}$$

§ 0.3.12 (necessity of $\to_{\mathrm{red}}$ and $\to_{\mathrm{nil}}$)**.** The purpose of $\to_{\mathrm{red}}$ together with $\to_{\mathrm{nil}}$ is to prevent unbounded growth of binding groups during unfolding. Consider for instance the outermost rewrite sequence on the term $\mathsf{let}\ f = \mathsf{let}\ g = f\ g\ \mathsf{in}\ g\ \mathsf{in}\ f$ shown in fig. 0.1, where after the fourth rewriting step g' (the left one) becomes unreachable and could be removed by a $\to_{\mathrm{red}}$-step.

While restricting the size of binding groups during unfolding is a sensible constraint on the unfolding process, it is not strictly necessary to define the unfolding of a $\lambda_{\mathsf{letrec}}$-term. Alternatively, we could employ a rule as follows:

$$\mathsf{let}\ f_1 = L_1 \ldots f_n = L_n\ \mathsf{in}\ L \to_{\mathrm{free}} L \quad (\text{if } f_1, \ldots, f_n \text{ do not occur in } L)$$

Note that a $\to_{\mathrm{free}}$-step can be simulated by a $\to_{\mathrm{red}}$-step followed by a $\to_{\mathrm{nil}}$-step. We will, however, at a later point embed the unfolding rules into other rewriting systems of which we wish to perform unfolding in a lazy way such that the number of derivable subterms is bounded. Using the $\to_{\mathrm{free}}$-rule instead of $\to_{\mathrm{red}}$ and $\to_{\mathrm{nil}}$ the size of the bindings groups in fig. 0.1 keeps growing and we obtain an infinite number of different subterms.

§ 0.3.13 (informal notation)**.** We will be using the informal notation as above throughout most of the thesis. For reasoning however, we lean on the theory of higher-order rewriting. In this way we can avoid the ado of an explicit substitution calculus, which would be required for sound reasoning in a first-order formulation.

$$
\begin{array}{ll}
 & \mathsf{let}\ f = \mathsf{let}\ g = f\ g\ \mathsf{in}\ g\ \mathsf{in}\ f \\
\to_{\mathrm{rec}} & \mathsf{let}\ f = \mathsf{let}\ g = f\ g\ \mathsf{in}\ g\ \mathsf{in}\ \mathsf{let}\ g' = f\ g'\ \mathsf{in}\ g' \\
\to_{\mathrm{letrec}} & \mathsf{let}\ \begin{array}{l} f = \mathsf{let}\ g = f\ g\ \mathsf{in}\ g \\ g' = f\ g' \end{array}\ \mathsf{in}\ g' \\
\to_{\mathrm{rec}} & \mathsf{let}\ \begin{array}{l} f = \mathsf{let}\ g = f\ g\ \mathsf{in}\ g \\ g' = f\ g' \end{array}\ \mathsf{in}\ f\ g' \\
\to_{@} & \left(\mathsf{let}\ \begin{array}{l} f = \mathsf{let}\ g = f\ g\ \mathsf{in}\ g \\ g' = f\ g' \end{array}\ \mathsf{in}\ f\right)\left(\mathsf{let}\ \begin{array}{l} f = \mathsf{let}\ g = f\ g\ \mathsf{in}\ g \\ g' = f\ g' \end{array}\ \mathsf{in}\ g'\right) \\
\to_{\mathrm{rec}} & \left(\mathsf{let}\ \begin{array}{l} f = \mathsf{let}\ g = f\ g\ \mathsf{in}\ g \\ g' = f\ g' \end{array}\ \mathsf{in}\ \mathsf{let}\ g'' = f\ g''\ \mathsf{in}\ g''\right) \\
 & \quad \left(\mathsf{let}\ \begin{array}{l} f = \mathsf{let}\ g = f\ g\ \mathsf{in}\ g \\ g' = f\ g' \end{array}\ \mathsf{in}\ g'\right) \\
\to_{\mathrm{letrec}} & \left(\mathsf{let}\ \begin{array}{l} f = \mathsf{let}\ g = f\ g\ \mathsf{in}\ g \\ g' = f\ g' \\ g'' = f\ g'' \end{array}\ \mathsf{in}\ g''\right)\left(\mathsf{let}\ \begin{array}{l} f = \mathsf{let}\ g = f\ g\ \mathsf{in}\ g \\ g' = f\ g' \end{array}\ \mathsf{in}\ g'\right)
\end{array}
$$

Figure 0.1. Unbounded growth of binding groups indicated by the initial segment of an infinite $\to_{\triangledown}$-rewrite-sequence without $\to_{\mathrm{red}}$-steps.

0.4 Preliminaries

Notation 0.4.1 ($\mathbb{N}$, natural numbers)**.** By $\mathbb{N}$ we denote the natural numbers including zero. We let $\mathbb{N} = \{0, 1, \ldots\}$.

Notation 0.4.2 (functions, domain, image)**.** For a total function $f : A \to B$ we denote by $\mathrm{dom}(f)$ the *domain* A, and by $\mathrm{im}(f)$ the *image* B of f.

For a partial function $f : A \rightharpoonup B$, and $a \in A$ we denote by $f(a)\!\downarrow$ that f is defined for a. The *domain* of f is the set $\mathrm{dom}(f) := \{a \in A \mid f(a)\!\downarrow\}$.

Notation 0.4.3. We denote by $f|_D$ the restriction of function f to domain D.

Notation 0.4.4 (rewrite relations)**.** Let $\to \subseteq A \times A$ be a rewrite relation. We denote by $\twoheadrightarrow$ the *many-step* rewrite relation induced by $\to$, by which we mean the reflexive and transitive closure of $\to$. By $\to^+$ we denote the *one-or-more-step* rewrite relation of $\to$, the transitive closure of $\to$. By $\to^=$ we mean the *zero-or-one-step* rewrite relation of $\to$, the reflexive closure of $\to$. By $\twoheadrightarrow$ we denote the

infinite rewrite relation of finitely of infinitely many $\to$-steps (see § 1.3.25 for more details). By a normal form of $\to$ we mean an $a \in A$ such that there is no $a' \in A$ with $a \to a'$. By $\to^!$ we mean the *reduction to normal form* rewrite relation induced by $\to$. It is equivalent to the restriction of $\twoheadrightarrow$ to a relation with the normal forms of $\to$ as codomain: $\to^! = \{\langle a, a' \rangle \mid a \twoheadrightarrow a', a' \text{ normal form of } \to\}$.

0.5 λ-terms and $\lambda_{\mathsf{letrec}}$-terms – CRS formalisation

§ 0.5.1 (Combinatory Reduction Systems)**.** Many of the formalisations we introduce are based on the framework of Combinatory Reduction Systems (CRSs) [36], [37] [51, Section 11.3], and, in particular, on infinitary Combinatory Reduction Systems (iCRSs) [31]. CRSs are a higher-order term rewriting framework tailor-made for formalising and manipulating expressions in higher-order languages (i.e. languages with binding constructs like λ-abstractions and function bindings). They provide a sound basis for defining our language and for reasoning with letrec-expressions. By formalising a system of unfolding rules as a CRS we conveniently externalise issues like name capturing and α-renaming, which otherwise would have to be handled by a calculus of explicit substitution. Also, we can lean on the rewriting theory of CRSs for the proofs.

Remark 0.5.2 (infinitary rewriting)**.** We rely on CRSs (and not for instance Higher-Order Rewriting Systems (HRSs) [51]) as a rewriting framework since to date infinitary rewriting theory has only been developed for CRSs [34, 35, 31].

§ 0.5.3 (the 'calculi' $\boldsymbol{\lambda}$, $\boldsymbol{\lambda}_\bullet$, and $\boldsymbol{\lambda}_{\mathsf{letrec}}$)**.** We will use the symbols $\boldsymbol{\lambda}$, $\boldsymbol{\lambda}_\bullet$, and $\boldsymbol{\lambda}_{\mathsf{letrec}}$ to refer to the λ-calculus, the λ-calculus with black holes, and the $\lambda_{\mathsf{letrec}}$-calculus with letrec. However, we consider the former two calculi not to have any rewriting rules at all, since we concern ourselves only with unfolding, not with β-reduction.

For formulating the rules from section 0.3 as a CRS, we provide CRS signatures for $\boldsymbol{\lambda}$, $\boldsymbol{\lambda}_\bullet$ and $\boldsymbol{\lambda}_{\mathsf{letrec}}$.

Definition 0.5.4 (CRS signatures for $\boldsymbol{\lambda}$, $\boldsymbol{\lambda}_\bullet$, and $\boldsymbol{\lambda}_{\mathsf{letrec}}$)**.** The CRS signature for $\boldsymbol{\lambda}$ consists of the set $\Sigma^\lambda = \{\mathsf{app}, \mathsf{abs}\}$ where app is a binary ($\mathrm{ar}(\mathsf{app}) = 2$) and abs is a unary ($\mathrm{ar}(\mathsf{abs}) = 1$) function symbol. The CRS signature for $\boldsymbol{\lambda}_\bullet$ is $\Sigma^\lambda_\bullet = \Sigma^\lambda \cup \{\bullet\}$ where $\bullet$ is a nullary ($\mathrm{ar}(\bullet) = 0$) function symbol. The CRS signature $\Sigma^\lambda_{\mathsf{letrec}}$ consists of the countably infinite set $\Sigma^\lambda_{\mathsf{letrec}} = \Sigma^\lambda_\bullet \cup \{\mathsf{let}\} \cup \{\mathsf{in}_n \mid n \in \mathbb{N}\}$ of function symbols, with $\mathrm{ar}(\mathsf{let}) = 1$ and $\mathrm{ar}(\mathsf{in}_n) = n + 1$ for all $n \in \mathbb{N}$.

Definition 0.5.5 (set of λ-terms)**.** By $\mathrm{Ter}^{\infty}(\boldsymbol{\lambda})$ we denote the set of closed iCRS terms (terms in an infinitary CRS [31]) over Σ^{λ}. Likewise, we denote by $\mathrm{Ter}^{\infty}(\boldsymbol{\lambda}_{\bullet})$ the set of closed iCRS terms over $\Sigma^{\lambda}_{\bullet}$. Note that the set includes finite as well as infinite terms, thus whenever we speak of λ-terms we refer to λ-terms that are either finite or infinite.

We will look more closely at iCRS-terms in § 1.3.22

Definition 0.5.6 (set of $\lambda_{\mathsf{letrec}}$-terms)**.** By $\mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$ we denote the set of closed CRS terms over $\Sigma^{\lambda}_{\mathsf{letrec}}$, with the restrictions

- that there is no occurrence of the $\bullet$-symbol
- that let and in_n can only occur as patterns of the form

$$\mathsf{let}([f_1 \dots f_n]\mathsf{in}_n(\dots))$$

- and that otherwise a CRS abstraction can only occur directly beneath an abs-symbol.

Example 0.5.7 (fix id)**.** The naive version of fix applied to the identity function as in example 0.1.11 in CRS notation:

$$\mathsf{app}(\mathsf{let}([\mathrm{fix}]\mathsf{in}_1(\mathsf{abs}([f]\mathsf{app}(f, \mathsf{app}(\mathrm{fix}, f))), \mathrm{fix})), \mathsf{abs}([x]x))$$

The unfolding of fix in CRS notation:

$$\mathsf{abs}([f]\mathsf{app}(f, \mathsf{app}(\mathsf{abs}([f]\mathsf{app}(f, \mathsf{app}(\dots, f))), f)))$$

Example 0.5.8 (fix)**.** The (not so naively implemented) fix-function from example 0.3.6 in CRS notation:

$$\mathsf{abs}([f](\mathsf{let}([r]\mathsf{in}_1(\mathsf{app}(f, r), r))))$$

0.6 Unfolding Rules – CRS formalisation

Here we give a CRS formalisation of the rules for unfolding $\lambda_{\mathsf{letrec}}$-terms, corresponding to unfolding as described informally in section 0.3.

Definition 0.6.1 (CRSs $\boldsymbol{R}_{\triangledown}$ and $\boldsymbol{R}_{\triangledown_{\bullet}}$ for unfolding $\lambda_{\mathsf{letrec}}$-terms)**.** $\boldsymbol{R}_{\triangledown}$ and $\boldsymbol{R}_{\triangledown_{\bullet}}$ for unfolding $\lambda_{\mathsf{letrec}}$-terms are CRSs over the signature $\Sigma^{\lambda}_{\mathsf{letrec}}$. The rules of $\boldsymbol{R}_{\triangledown_{\bullet}}$ consist of all the rule schemes below, while in $\boldsymbol{R}_{\triangledown}$ the last two rules schemes

are excluded ($\varrho_{\triangledown}^{\text{tighten}}$ and $\varrho_{\triangledown}^{\bullet}$). We use vector notation to denote sequences of CRS abstractions ($\vec{f}$ instead of $f_1, \ldots, f_n$) and metaterms ($\vec{B}(\vec{f})$ instead of $B_1(\vec{f}), \ldots, B_n(\vec{f})$).

$$\varrho^{\lambda}_{\triangledown}:\ \mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(\vec{B}(\vec{f}),\mathsf{abs}\big([x]\,M(\vec{f},x)\big)\big)\big)$$
$$\to \mathsf{abs}\big([x]\,\mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(\vec{B}(\vec{f}),M(\vec{f},x)\big)\big)\big)$$

$$\varrho^{@}_{\triangledown}:\ \mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(\vec{B}(\vec{f}),\mathsf{app}\big(M(\vec{f}),\,N(\vec{f})\big)\big)\big)$$
$$\to \mathsf{app}\Big(\begin{array}{c}\mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(\vec{B}(\vec{f}),M(\vec{f})\big)\big),\\ \mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(\vec{B}(\vec{f}),N(\vec{f})\big)\big)\end{array}\Big)$$

$$\varrho^{\mathrm{letrec}}_{\triangledown}:\ \mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(\vec{B}(\vec{f}),\mathsf{let}\big([\vec{g}]\,\mathsf{in}_m\big(\vec{C}(\vec{f},\vec{g}),M(\vec{f},\vec{g})\big)\big)\big)\big)$$
$$\to \mathsf{let}\big([\vec{f}\vec{g}]\,\mathsf{in}_{n+m}\big(\vec{B}(\vec{f}),\vec{C}(\vec{f},\vec{g}),M(\vec{f},\vec{g})\big)\big)$$

$$\varrho^{\mathrm{rec}}_{\triangledown}:\ \mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(\vec{B}(\vec{f}),f_i\big)\big)\to\mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(\vec{B}(\vec{f}),B_i(\vec{f})\big)\big)$$

$$\varrho^{\mathrm{nil}}_{\triangledown}:\ \mathsf{let}(\mathsf{in}_0(M))\to M$$

$$\varrho^{\mathrm{red}}_{\triangledown}:\ \mathsf{let}\Big([\vec{f}]\,\mathsf{in}_n\Big(\fatsemi_{i\in\{1,\ldots,n\}} B_i(\vec{f}_i)\ ,\quad M(\vec{f}')\Big)\Big)$$
$$\to \mathsf{let}\Big([\vec{f}']\,\mathsf{in}_{|I|}\Big(\fatsemi_{i\in I} B_i(\vec{f}')\ ,\quad M(\vec{f}')\Big)\Big)$$
for some $I\subset\{1,\ldots,n\}$
with $\vec{f}'=\fatsemi_{i\in I} f_i$ and $\vec{f}_i=\begin{cases}\vec{f}' & i\in I\\ \vec{f} & i\notin I\end{cases}$

$$\varrho^{\mathrm{tighten}}_{\triangledown}:\ \mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(\vec{B}(\vec{f}),M(\vec{f})\big)\big)\text{ where }B_i(\vec{f})=f_j$$
$$\to \mathsf{let}\Big([\vec{g}]\,\mathsf{in}_{n-1}\Big(\begin{array}{c}B_1(\vec{g}'),\ldots,B_{i-1}(\vec{g}'),\\ B_{i+1}(\vec{g}'),\ldots,B_n(\vec{g}'),M(\vec{g}')\end{array}\Big)\Big)$$
for some $i,j\in\{1,\ldots,n\}$ with $i\neq j$
and $\vec{g}'=\langle g_1,\ldots,g_{i-1},g_k,g_{i+1},\ldots,g_{n-1}\rangle$
where $k=j$ if $j<i$ and $k=j-1$ if $j>i$

$$\varrho^{\bullet}_{\triangledown}:\ \mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(\vec{B}(\vec{f}),M(\vec{f})\big)\big)\text{ where }B_i(\vec{f})=f_i$$
$$\to \mathsf{let}\Big([\vec{g}]\,\mathsf{in}_{n-1}\Big(\begin{array}{c}B_1(\vec{g}'),\ldots,B_{i-1}(\vec{g}'),\\ B_{i+1}(\vec{g}'),\ldots,B_n(\vec{g}'),M(\vec{g}')\end{array}\Big)\Big)$$
for some $i\in\{1,\ldots,n\}$
and $\vec{g}'=\langle g_1,\ldots,g_{i-1},\bullet,g_{i+1},\ldots,g_{n-1}\rangle$

The first four rule schemes require little explanation. The rules are generated from having n and m range over $\mathbb{N}$. The length of the vectors $\vec{f}$, $\vec{g}$, $\vec{B}$, and $\vec{C}$ are stipulated by the subscript of the let-symbol in accordance to definition 0.5.6. $\varrho^{\mathrm{nil}}_{\triangledown}$ is actually a rule, not a rule scheme.

The formulation of $\varrho^{\mathrm{red}}_{\triangledown}$ uses, as an abbreviation, an ad-hoc list-builder notation, which works just as customary mathematical notation for, say, the union by indexing over an ordered set. Compare:

$$\bigcup_{i\in\{i_1,\ldots,i_n\}} A(i) = A(i_1)\cup\ldots\cup A(i_n) \qquad \text{where } i_1<\ldots<i_n$$
$$\fatsemi_{i\in\{i_1,\ldots,i_n\}} A(i) = A(i_1),\ldots,A(i_n)$$

Moreover, in the rule scheme $\varrho^{\mathrm{red}}_{\triangledown}$ not only does n range over $\mathbb{N}$; also the index set I ranges over all subsets of $\{1,\ldots,n\}$. The purpose of the rule scheme is to remove all bindings that are not *required*. A binding is considered required if it is used directly or indirectly by M. I is (by formulation of the rule scheme) a superset of all required bindings, or to be more precise: for a given term, only if I is chosen as a superset of the required bindings, the rule scheme yields an CRS rule applicable to the term. From the set of 'valid' choices we consider the single minimal choice for I in order to remove *all* unrequired bindings at once. An implementation of this rule scheme would entail a reachability analysis, which is implicit here.

Example 0.6.2 ($\varrho^{\mathrm{red}}_{\triangledown}$)**.** To understand the rule scheme $\varrho^{\mathrm{red}}_{\triangledown}$, consider the term $L = \mathsf{let}\ f_1 = f_2,\ f_2 = f_1\ \mathsf{in}\ f_1$ or in CRS notation $L = \mathsf{let}([f_1 f_2]\mathsf{in}_2(f_2, f_1, f_1))$. Considering only the four possibilities we have for I, $\varrho^{\mathrm{red}}_{\triangledown}$ induces the following four CRS rules:

$$\begin{aligned}
\{\}:\ & \mathsf{let}([f_1 f_2]\mathsf{in}_2(B_1(f_1,f_2),B_2(f_1,f_2),M())) \\
& \quad\to\quad \mathsf{let}(\mathsf{in}_0(M())) \\
\{1\}:\ & \mathsf{let}([f_1 f_2]\mathsf{in}_2(B_1(f_1),B_2(f_1,f_2),M(f_1))) \\
& \quad\to\quad \mathsf{let}([f_1]\mathsf{in}_1(B_1(f_1),M(f_1))) \\
\{2\}:\ & \mathsf{let}([f_1 f_2]\mathsf{in}_2(B_1(f_1,f_2),B_2(f_2),M(f_2))) \\
& \quad\to\quad \mathsf{let}([f_2]\mathsf{in}_1(B_2(f_2),M(f_2))) \\
\{1,2\}:\ & \mathsf{let}([f_1 f_2]\mathsf{in}_2(B_1(f_1,f_2),B_2(f_1,f_2),M(f_1,f_2))) \\
& \quad\to\quad \mathsf{let}([f_1,f_2]\mathsf{in}_2(B_2(f_1,f_2),M(f_1,f_2)))
\end{aligned}$$

The first and the third rule are not applicable to L, since M has an occurrence of f_1. The rules induced by $I = \{1\}$ and $I = \{2\}$ are not applicable to L, because B_1 has an occurrence of f_2 and B_2 has an occurrence of f_1. All the bindings here are used by M, therefore the only applicable rule is the last one, which does not alter L at all.

Now let us consider the term $L' = \mathsf{let}\ f_1 = f_2,\ f_2 = f_1\ \mathsf{in}\ x$ or in CRS notation $L' = \mathsf{let}([f_1 f_2]\mathsf{in}_2(f_2, f_1, x))$. As before the rules induced by $I = \{1\}$ and $I = \{2\}$ are not applicable. However this time not only the last but also the first rule is applicable, indicating that none of the bindings are used by M and thus all of them can be removed in one $\to_{\mathrm{red}}$-step.

Definition 0.6.3 (garbage free)**.** We call a $\lambda_{\mathsf{letrec}}$-term that is a normal form w.r.t. to the rules $\varrho^{\mathrm{red}}_{\triangledown}$ and $\varrho^{\mathrm{nil}}_{\triangledown}$ *garbage free.*

Notation 0.6.4 (ϱ^{ρ}, ϱ^{ρ}_{s})**.** Above we used the notation ϱ^{ρ}_{s} to denote a rule named ρ that belongs to the rewriting system s. We will omit the s and write ϱ^{ρ} to denote a rule ρ if it is unambiguous to which rewriting system ρ belongs to.

Notation 0.6.5 (rewriting relations for $\boldsymbol{R}_{\triangledown}$ and $\boldsymbol{R}_{\triangledown_\bullet}$)**.** We write $\to_{\triangledown}$ ($\to_{\triangledown_\bullet}$) for the rewrite relation induced by $\boldsymbol{R}_{\triangledown}$ ($\boldsymbol{R}_{\triangledown_\bullet}$). And by $\to_{\lambda}$, $\to_{@}$, $\to_{\mathrm{letrec}}$, $\to_{\mathrm{rec}}$, $\to_{\mathrm{nil}}$, $\to_{\mathrm{red}}$, $\to_{\mathrm{tighten}}$, and $\to_{\bullet}$, we denote the rewrite relations of both the CRSs that are induced by the rules ϱ^{λ}, $\varrho^{@}$, $\varrho^{\mathrm{letrec}}$, ϱ^{rec}, ϱ^{nil}, ϱ^{red}, $\varrho^{\mathrm{tighten}}$, and $\varrho^{\bullet}$, respectively.

Remark 0.6.6 (alternative formalisation with permutations)**.** Note that in the rule patterns in $\boldsymbol{R}_{\triangledown}$ we have to ensure that the function binding we want to refer to may be at an arbitrary position among the function bindings of a let-expression. An alternative approach would be to adding a permutation rule by which two adjacent function bindings can be swapped. Then the remaining rules which refer to a function binding could be written such that they always make use of the first function binding of the let-expression. This approach is for instance used in [17].

§ 0.6.7 (informal notation)**.** As mentioned in § 0.3.13 when studying examples we will mostly rely on the easier-to-read informal notation from section 0.3 instead of the more cumbersome CRS notation.

0.7 Unfolding Semantics of $\lambda_{\mathsf{letrec}}$

Based on the CRSs $\boldsymbol{R}_{\nabla}$ and $\boldsymbol{R}_{\nabla_\bullet}$ we define the unfolding semantics of $\lambda_{\mathsf{letrec}}$-terms as its infinite unique normal form. The well-definedness of the functions below is witnessed by theorem 0.7.2 below.

Definition 0.7.1 (partial and total unfolding semantics)**.** $\boldsymbol{R}_{\nabla}$ induces the unfolding semantics

$$\begin{aligned}[\![\cdot]\!]_{\boldsymbol{\lambda}} : \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}}) &\rightharpoonup \mathrm{Ter}^{\infty}(\boldsymbol{\lambda}) \\ L &\mapsto M \quad \text{if } L \twoheadrightarrow^{!}_{\nabla} M\end{aligned}$$

which is partial, because of meaningless bindings. (see § 0.3.7)

$\boldsymbol{R}_{\nabla_\bullet}$ induces the unfolding semantics

$$\begin{aligned}[\![\cdot]\!]_{\boldsymbol{\lambda}_\bullet} : \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}}) &\to \mathrm{Ter}^{\infty}(\boldsymbol{\lambda}_\bullet) \\ L &\mapsto M \quad \text{if } L \twoheadrightarrow^{!}_{\nabla_\bullet} M\end{aligned}$$

which is total, as it maps meaningless terms to $\bullet$.

Theorem 0.7.2. $[\![\cdot]\!]_{\boldsymbol{\lambda}}$ and $[\![\cdot]\!]_{\boldsymbol{\lambda}_\bullet}$ are well-defined functions.

Proof. Well-definedness of these mappings is guaranteed by the following properties of $\boldsymbol{R}_{\nabla}$ and $\boldsymbol{R}_{\nabla_\bullet}$:

Lemma 0.7.6 (infinite normalisation): the existence of a normal form (only required for $\boldsymbol{R}_{\nabla_\bullet}$)

Lemma 0.7.5 (well-formedness of the normal forms): normal forms do indeed adhere to the signatures Σ^{λ}, and $\Sigma^{\lambda}_{\bullet}$, respectively.

Lemma 0.7.7 (uniqueness of normal forms)

□

Proposition 0.7.3 ($[\![\cdot]\!]_{\boldsymbol{\lambda}}$ is a specialisation $[\![\cdot]\!]_{\boldsymbol{\lambda}_\bullet}$)**.**

$$\forall L \in \mathrm{dom}([\![\cdot]\!]_{\boldsymbol{\lambda}})\ [\![L]\!]_{\boldsymbol{\lambda}} = [\![L]\!]_{\boldsymbol{\lambda}_\bullet}$$

Proof. For $\boldsymbol{R}_{\nabla}$ as well as $\boldsymbol{R}_{\nabla_\bullet}$ a $\lambda_{\mathsf{letrec}}$-term is a normal form if and only if it is a let-expression. Therefore every normal form of $[\![\cdot]\!]_{\boldsymbol{\lambda}}$ is a normal form of $[\![\cdot]\!]_{\boldsymbol{\lambda}_\bullet}$. Reachability of the normal forms is guaranteed by the fact that the rules of $\boldsymbol{R}_{\nabla_\bullet}$ are a superset of the rules of $\boldsymbol{R}_{\nabla}$. □

Terminology 0.7.4. We say that a $\lambda_{\mathsf{letrec}}$-term L *unfolds to* M if either $[\![L]\!]_{\boldsymbol{\lambda}} = M$ or $[\![L]\!]_{\boldsymbol{\lambda}_\bullet} = M$. Which is meant, should be clear from the context.

We finish the section with the lemmas required for theorem 0.7.2.

Lemma 0.7.5 (well-formedness of the normal forms)**.** $\boldsymbol{R}_{\nabla}$ ($\boldsymbol{R}_{\nabla_\bullet}$) has λ-terms ($\lambda_\bullet$-terms) as normal forms.

Proof sketch. The names of the first four rules are chosen to reflect the kind of term contained by the body of the letrec-expression, which helps to see that the rules are complete in the sense that every let-expression is a redex, thus normal forms do not contain let-expressions. Terms over $\Sigma^{\lambda}_{\mathsf{letrec}}$ without let-expressions are $\lambda_\bullet$-terms. The arguments also holds analogously for $\boldsymbol{R}_{\nabla}$, which contains no rule to give rise to a $\bullet$-symbol. □

Lemma 0.7.6 (infinite normalisation of $\boldsymbol{R}_{\nabla_\bullet}$)**.** Every $\lambda_{\mathsf{letrec}}$-term is either weakly normalising w.r.t. $\to_{\nabla_\bullet}$ or admits a strongly convergent outermost-fair $\to_{\nabla_\bullet}$-rewriting-sequence.

For definitions of 'outermost fair' and 'strongly convergent' we refer to [35].

Proof sketch. We consider outermost-fair $\to_{\nabla_\bullet}$-rewrite-sequences on a $\lambda_{\mathsf{letrec}}$-term L, in which ϱ^{red}, ϱ^{nil}, $\varrho^{\mathrm{tighten}}$, and $\varrho^{\bullet}$ are applied eagerly. We show that every such sequence τ is either finite, or that otherwise its rewrite activity tends to infinity. We argue by contradiction: We assume τ performs infinitely many rewriting steps on position p. Then there is an infinite subsequence ξ of τ which contracts only redexes at p. We show that ξ cannot exist.

First note that a $\lambda_{\mathsf{letrec}}$-term is a $\to_{\nabla_\bullet}$-redex if and only if it is a let-expression. Also an outermost $\to_{\nabla_\bullet}$-rewriting-sequence cannot create redexes upwards. Therefore, for ξ to be infinite all terms in ξ must have a let-expression at p.

ξ can contain neither $\to_{\lambda}$-steps nor $\to_{@}$-steps, since they would generate a function symbol at p (and thus yield a term that is not a let-expression).

For the rest of the proof we argue using the concept of letrec-*depth*, the number of let-symbols occurring under p.

ξ cannot contain an infinite number of $\to_{\mathrm{nil}}$-steps because it reduces letrec-depth and no other rule increases letrec-depth.

Thus, ξ must have an infinite suffix π which contains no $\to_{\lambda}$-, $\to_{@}$-, and $\to_{\mathrm{nil}}$-steps, but only steps due to $\to_{\mathrm{letrec}}$, $\to_{\mathrm{rec}}$, $\to_{\mathrm{red}}$, $\to_{\mathrm{tighten}}$, and $\to_{\bullet}$.

For π to be infinite it must contain infinitely many $\to_{\mathrm{rec}}$-steps since the remaining four rules are only finitely often applicable, because they manipulate or remove bindings (which can happen only once per binding), or in the case of $\to_{\mathrm{letrec}}$ decrease the letrec-depth.

We will now go on to show that π cannot have infinitely many $\to_{\mathrm{rec}}$-steps, if $\to_{\mathrm{tighten}}$-steps and $\to_{\bullet}$-steps take precedence. Every binding that is employed by the $\to_{\mathrm{rec}}$-step can only be of the form $f = g$ or $f = \mathsf{let} \ldots \mathsf{in} \ldots$, otherwise a $\to_{\lambda}$-step or a $\to_{@}$-step would ensue, which we already excluded. A term with only these two forms of bindings can be reduced to a term with letrec-depth 1 with only bindings of the form $f = g$ by applications of $\varrho^{\mathrm{tighten}}$, ϱ^{red}, ϱ^{nil}, ϱ^{rec}, $\varrho^{\mathrm{letrec}}$. If the bindings are cyclic, the π terminates with an application of $\varrho^{\bullet}$, otherwise the resulting term is a free variable. □

Lemma 0.7.7 (uniqueness of normal forms)**.** If for a term in $\mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$ a normal form exists with respect to $\boldsymbol{R}_{\nabla}$ or $\boldsymbol{R}_{\nabla_{\bullet}}$, it is unique.

Proof. This follows from finitary confluence (proposition 0.7.8) □

Proposition 0.7.8. $\boldsymbol{R}_{\nabla}$ and $\boldsymbol{R}_{\nabla_{\bullet}}$ are confluent.

Proof. Unique infinite normalisation of $\boldsymbol{R}_{\nabla}$ and $\boldsymbol{R}_{\nabla_{\bullet}}$ follows from finitary confluence of $\boldsymbol{R}_{\nabla}$ and $\boldsymbol{R}_{\nabla}$. In previous work [18] we proved confluence for a CRS which is very similar to $\boldsymbol{R}_{\nabla}$. In appendix C an adapted version of that proof is provided, also extended by the rules $\varrho^{\mathrm{tighten}}$ and $\varrho^{\bullet}$. The proof is based on decreasing diagrams [51, Section 14.2] and involves a comprehensive critical-pair analysis. □

Remark 0.7.9. Note that this confluence result concerns a rewriting system only for unfolding $\lambda_{\mathsf{letrec}}$-terms, and therefore does not conflict with non-confluence observations concerning versions of cyclic λ-calculi which include unfolding rules as well as β-reduction [3].

§ 0.7.10 ($\mathrm{dom}(\llbracket\cdot\rrbracket_{\boldsymbol{\lambda}})$)**.** Finally we are going to characterise the domain of $\llbracket\cdot\rrbracket_{\boldsymbol{\lambda}}$, thus identify those $\lambda_{\mathsf{letrec}}$-terms that are not meaningless but unfold to some λ-term. Meaningless $\lambda_{\mathsf{letrec}}$-terms are unproductive in the sense that during all outermost-fair $\boldsymbol{R}_{\nabla}$-rewrite-sequences the production of symbols stagnates due to an unproductive cycle.

Lemma 0.7.11. Let L be a $\lambda_{\mathsf{letrec}}$-term be a let-expression. Then exactly one of the following statements hold:

- All maximal outermost-fair $\boldsymbol{R}_{\triangledown}$-rewriting-sequences on L solely contain let-expressions.
- All maximal outermost-fair $\boldsymbol{R}_{\triangledown}$-rewriting-sequences on L contain finitely many let-expressions.

Definition 0.7.12 ($\boldsymbol{R}_{\triangledown}$-productivity)**.** We say that a $\lambda_{\mathsf{letrec}}$-term L is $\boldsymbol{R}_{\triangledown}$-*productive* if the following statement holds:

- L does not have a $\to_{\triangledown}$-reduct that is the source of an infinite $\to_{\triangledown}$-rewrite-sequence consisting exclusively of outermost steps with respect to $\to_{\mathrm{rec}}$, $\to_{\mathrm{nil}}$, $\to_{\mathrm{letrec}}$, or $\to_{\mathrm{red}}$.

Lemma 0.7.13. For every $\lambda_{\mathsf{letrec}}$-terms L the following statements are equivalent:

(i) $L \to_{\triangledown}^{\omega} M$ for some infinite λ-term M.

(ii) L is $\boldsymbol{R}_{\triangledown}$-productive.

(iii) Every maximal outermost-fair $\to_{\triangledown}$-rewrite-sequence on L is strongly convergent.

Proof. (ii) $\Rightarrow$ (iii), because if L is $\boldsymbol{R}_{\triangledown}$-productive then every outermost occurrence of a letrec in every $\to_{\triangledown}$-reduct will be eventually pushed down to a higher position by either a $\to_{\lambda}$-step or a $\to_{@}$-step of any maximal outermost-fair $\to_{\triangledown}$-sequence. Since only let-expressions are $\to_{\triangledown}$-redexes any maximal outermost-fair rewrite sequence starting from L converges to an infinite normal form. (i) follows directly from (iii). (i) $\Rightarrow$ (ii) follows from lemma 0.7.11 by contradiction. If L is not $\boldsymbol{R}_{\triangledown}$-productive then it has by definition a $\to_{\triangledown}$-reduct with at least one occurrence of a letrec which cannot be pushed further down by any outermost application of any $\boldsymbol{R}_{\triangledown}$-rule. By lemma 0.7.11 the same holds for every other maximal outermost-fair rewrite sequence. Therefore L cannot unfold to a λ-term M because M may not contain any letrecs. □

Proposition 0.7.14 ($\mathrm{dom}(\llbracket\cdot\rrbracket_{\boldsymbol{\lambda}})$)**.** A $\lambda_{\mathsf{letrec}}$-term is in $\mathrm{dom}(\llbracket\cdot\rrbracket_{\boldsymbol{\lambda}})$ if and only if it admits no cyclic $\boldsymbol{R}_{\triangledown}$-rewriting-sequence without $\varrho_{\triangledown}^{@}$-steps and $\varrho_{\triangledown}^{\lambda}$-steps.

Proof. "⇒" follows from confluence of $\boldsymbol{R}_{\triangledown}$ (proposition 0.7.8). For "⇐" we assume that L is not infinitarily normalising in $\boldsymbol{R}_{\triangledown}$. That means that any fair strategy rewrites M eventually to a reduct with root-active subterm O. The infinite rewriting sequence that only applies rules to the root of O can never contain any $\varrho^{@}_{\triangledown}$-steps or $\varrho^{\lambda}_{\triangledown}$-steps because the reduct would not be a redex. The cyclicity of this rewriting sequence follows from proposition 1.6.22. □

§ 0.7.15 (a single-rule unfolding system in the likeness of μ-unfolding)**.** In appendix D we present an alternative rewriting system for unfolding $\lambda_{\mathsf{letrec}}$-terms. It has only a single rule, like the canonical rewriting system for unfolding μ-terms.

§ 0.7.16 (outlook)**.** Now that we have established what we mean by unfolding, we can phrase the problems we tackle in the three following chapters of this thesis:

chapter 1: Here we characterise the set of λ-terms that can be expressed finitely as $\lambda_{\mathsf{letrec}}$-terms. To this end we introduce rewriting systems for deconstructing λ-terms. These rewriting systems induce a notion of regularity on a decomposed λ-term: if the set of its components is finite it is regular.

chapter 2: From the decompositions systems arise graphical representations for λ-terms in a natural way (essentially the reduction graph of a term w.r.t. a decomposition system). We study these graph representations in order to use them to reason about their respective λ-terms.

chapter 3: Here we relate these graph formalisms back to $\boldsymbol{\lambda}_{\mathsf{letrec}}$ and unfolding, by which we obtain concrete practical methods to analyse and manipulate $\lambda_{\mathsf{letrec}}$-terms.

Chapter 1

Expressibility in $\lambda_{\mathsf{letrec}}$

1.1 Overview

§ 1.1.0 (teaser)**.** Why cannot all infinite $\lambda_{\mathsf{letrec}}$-terms with a simple repetitive structure like $\lambda a.\, \lambda b.\, (\lambda a.\, (\lambda b. \ldots .\, a)\; b)\; a$ be expressed in $\boldsymbol{\lambda}_{\mathsf{letrec}}$?

§ 1.1.1 (subject matter)**.** In this chapter we study the relationship between finite terms in $\boldsymbol{\lambda}_{\mathsf{letrec}}$ and the λ-terms they express. We investigate λ-terms that are not unfoldings of any $\lambda_{\mathsf{letrec}}$-term and we consider the question: Which are the infinite lambda terms that are $\lambda_{\mathsf{letrec}}$-expressible in the sense that they can be obtained as infinite unfoldings of finite $\lambda_{\mathsf{letrec}}$-terms? Or in other words: how expressive is the language $\boldsymbol{\lambda}_{\mathsf{letrec}}$?

§ 1.1.2 (methods and formalisms)**.** We introduce a rewrite system for observing λ-terms through repeated experiments carried out at the head of the term, thereby decomposing it into 'generated subterms'. There are four sorts of decomposition steps: $\to_{\lambda}$-steps (decomposing a λ-abstraction), $\to_{@_0}$-steps and $\to_{@_1}$-steps (decomposing an application into its function and argument), and a scope-delimiting step. The scope-delimiting step comes in two variants, $\to_{\mathrm{del}}$ and $\to_{\mathsf{S}}$, defining two rewriting systems with rewrite relations $\to_{\mathrm{reg}}$ and $\to_{\mathrm{reg}^+}$. These rewrite relations each induce a notion of 'λ-transition-graph', a sort of graphical representation of λ-terms, which can be compared by means of bisimulation. We call a λ-term 'regular' ('strongly regular') if its set of $\to_{\mathrm{reg}}$-reachable ($\to_{\mathrm{reg}^+}$-reachable) generated subterms (and therefore its λ-transition-graph) is finite. Furthermore, we analyse the binding structure of λ-terms with the concept of 'binding–capturing chains'.

§ 1.1.3 (result)**.** Utilising these concepts, we answer the above question by providing two characterisations of $\lambda_{\mathsf{letrec}}$-expressibility. For all λ-terms M, the following statements are equivalent: (i): M is $\lambda_{\mathsf{letrec}}$-expressible; (ii): M is strongly regular; (iii): M is regular, and it only has finite binding–capturing chains.

1.2 Introduction

In chapter 0 we have established how $\lambda_{\mathsf{letrec}}$-terms serve as a finite representation of (potentially) infinite λ-terms. It is quite obvious, that not every infinite λ-term can be represented finitely, only those λ-terms come into consideration that have some kind of repetitive, or regular, structure. It turns out, however, that not even λ-terms with a regular syntax tree can always be expressed as the infinite unfolding of a term in $\boldsymbol{\lambda}_{\mathsf{letrec}}$ (see example 1.2.2 below).

Terminology 1.2.1 ($\lambda_{\mathsf{letrec}}$-expressible)**.** We say that a λ-term M is $\lambda_{\mathsf{letrec}}$-*expressible* if it has a representation as a (finite)[1] term L in $\boldsymbol{\lambda}_{\mathsf{letrec}}$:

$$\exists L \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}}) \quad M = [\![L]\!]_{\boldsymbol{\lambda}}$$

We then also say that L *expresses* M.

Note that λ-terms have infinitely many different representations as $\lambda_{\mathsf{letrec}}$-terms (see e.g. example 1.2.5 below).

Example 1.2.2 (not $\lambda_{\mathsf{letrec}}$-expressible)**.** Consider the infinite λ-term of the form $M = \lambda a.\, \lambda b.\, (\lambda c.\, (\lambda d. \ldots .\, c)\; b)\; a$ with syntax trees as shown in fig. 1.1. Even though it has a syntax tree with a regular structure, M is not $\lambda_{\mathsf{letrec}}$-expressible.

§ 1.2.3 (the relevance of scopes)**.** To understand why there is no $\lambda_{\mathsf{letrec}}$-term that unfolds to M, it helps to consider the scopes of the abstractions in M. Informally speaking, the scope of an abstraction denotes the minimal connected portion of the term which includes the abstraction itself as well as all occurrences of the bound variable that it binds. A precise definition is given later in definition 1.7.9. The scopes of M as shown in fig. 1.1 are infinitely entangled: the scope of λa reaches into the scope of λb, the scope of λb into the scope of λc, and so on. This trait of M suggests M cannot be the result of 'unrolling' a $\lambda_{\mathsf{letrec}}$-term L, since such a process would map L's scoping structure onto M in a regular manner, and result in a term which is 'tiled' into finite, non-overlapping scopes with

[1] Recall that $\lambda_{\mathsf{letrec}}$-terms are finite by definition 0.5.6.

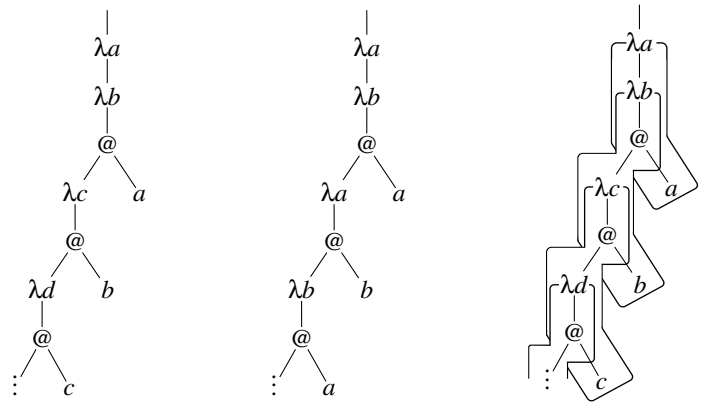

Figure 1.1. Two (α-equivalent) syntax trees of M from example 1.2.2. The second syntax tree is cleary regular. The syntax tree on the right is a version of the first syntax tree with scopes depicted as box-like structures.

bounded scope-nesting depth. This excludes, intuitively, the formation of the infinite entanglement of successively overlapping scopes that can be observed in M.

Example 1.2.4 ($\lambda_{\mathsf{letrec}}$-expressible)**.** Consider the infinite λ-term of the form $M = \lambda xy.\, M\ y\ x$. It is $\lambda_{\mathsf{letrec}}$-expressible as it arises as the unfolding of $\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f$ as witnessed by the following rewrite sequence:

$$
\begin{aligned}
& \mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f \\
\to_{\mathrm{rec}}\ & \mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ \lambda xy.\, f\ y\ x \\
\to_{\lambda}\ & \lambda x.\, \mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ \lambda y.\, f\ y\ x \\
\to_{\lambda}\ & \lambda xy.\, \mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f\ y\ x \\
\to_{@}\ & \lambda xy.\, (\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f\ y)\ (\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ x) \\
\to_{\mathrm{red}}\ & \lambda xy.\, (\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f\ y)\ (\mathsf{let}\ \ \mathsf{in}\ x) \\
\to_{\mathrm{nil}}\ & \lambda xy.\, (\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f\ y)\ x \\
\to_{@}\ & \lambda xy.\, (\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f)\ (\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ y)\ x \\
\to_{\mathrm{red}}\ & \lambda xy.\, (\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f)\ (\mathsf{let}\ \ \mathsf{in}\ y)\ x \\
\to_{\mathrm{nil}}\ & \lambda xy.\, (\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f)\ y\ x
\end{aligned}
$$

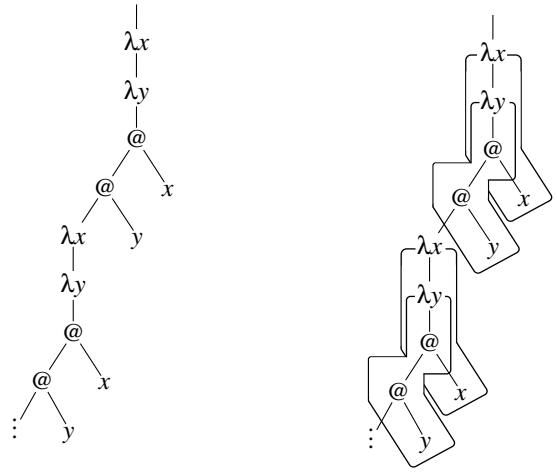

Figure 1.2. Syntax tree of M from example 1.2.4, and a version annotated with scopes.

$\rightarrow_{\mathrm{rec}} \ \cdots$

As shown in fig. 1.2 the syntax tree of M has some entanglement (the scopes of λx and λy do overlap) but the entanglement is finite as opposed to the entanglement in example 1.2.2.

Example 1.2.5 ($\lambda_{\mathsf{letrec}}$-expressible)**.** The infinite λ-term $\lambda f.\, f\ (f\ (f\ \ldots))$ from example 0.3.6 can be expressed by L as well as by P defined as follows:

$$L \coloneqq \lambda f.\, \mathsf{let}\ r = f\ r\ \mathsf{in}\ r \qquad P \coloneqq \lambda f.\, \mathsf{let}\ r = f\ (f\ r)\ \mathsf{in}\ r$$

It holds that $[\![L]\!]_{\boldsymbol{\lambda}} = \lambda f.\, f\ (f\ (f\ \ldots)) = [\![P]\!]_{\boldsymbol{\lambda}}$. See fig. 1.3 for the corresponding 'syntax graphs'. There is only one abstraction, so trivially there are no overlapping scopes.

§ 1.2.6 (regularity of first-order terms)**.** Over a first order signature, the simplest kind of infinite terms are those that are regular in the sense that they possess only a finite number of different subterms. They correspond to trees

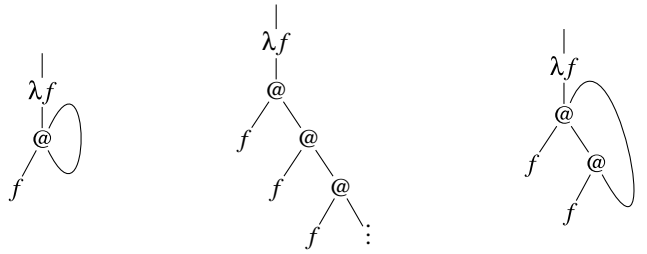

$\lambda f.\,\mathsf{let}\ r = f\ r\ \mathsf{in}\ r \twoheadrightarrow_{\triangledown} \quad \lambda f.\,f\ (f\ (f\ \ldots)) \quad \twoheadleftarrow_{\triangledown} \lambda f.\,\mathsf{let}\ r = f\ (f\ r)\ \mathsf{in}\ r$

Figure 1.3. The 'syntax graphs' of the $\lambda_{\mathsf{letrec}}$-terms from example 1.2.5 and the syntax tree of their unfolding.

over ranked alphabets that are regular [13]. Like regular trees, also regular terms can be expressed finitely by systems of recursion equations [13] or by 'rational expressions' [13, Definition 4.5.3], which correspond to μ-terms (see e.g. [15]). Hereby finite expressions denote infinite terms either via a mathematical definition (a fixed-point construction, or induction on paths) or as the limit of a rewrite sequence consisting of unfolding steps.

§ 1.2.7 (regularity of higher-order terms)**.** For higher-order terms such as λ-terms the concept of regularity is less clear-cut from the outset, due to the presence of variable bindings. Frequently, regularity has been used to denote the existence of a first-order representation with named variables that is regular (e.g. in [3, 1]). According to that definition M from example 1.2.2 would be regular: while the syntax tree on the left in fig. 1.1 of the term M contains infinitely many variables (and therefore is not regular), M has another α-equivalent regular syntax tree (on the right) that uses only two variable names. However, such a definition of regularity has two drawbacks: it does not (as desired) correspond to $\lambda_{\mathsf{letrec}}$-expressibility; and it relies on a property of a first-order representation. We will attempt to find a stronger, more direct definition of regularity for λ-terms that arises as an adaptation of the first-order notion of regularity.

§ 1.2.8 (subterms of higher-order terms)**.** The first-order notion of regularity relies on the finiteness of the number of subterms. When adapting this notion to higher-order terms, the problem is that there is no immediate, self-evident

notion of subterm. Should for instance y be considered a subterm of $\lambda x.\, x$ which is after all equivalent to $\lambda y.\, y$? In order to arrive at a viable notion of higher-order subterms we define a CRS for decomposing higher-order terms. The decomposition 'remembers' the abstractions encountered and 'saves' them as part of the such obtained components which we call 'generated subterms'.

§ 1.2.9 (prefixed subterms)**.** Viable notions of subterms for λ-terms in a higher-order formalisation require a stipulation on how to treat variable binding when descending into the body of a λ-abstraction. For this purpose we enrich the syntax of λ-terms with a parenthesised prefix of abstractions (similar to a proof system for weak μ-equality in [15, Figure 12]). An expression $(\lambda x_1 \dots x_n)\ M$ represents a partially decomposed λ-term: the body M typically contains free occurrences of variables which in the original λ-term were bound by λ-abstractions that have been split off by decomposition steps. The role of such abstractions has then been taken over by abstractions in the prefix $(\lambda x_1 \dots x_n)$. In this way expressions with abstraction prefixes can be kept closed under decomposition steps.

§ 1.2.10 (decomposition CRSs)**.** On these prefixed λ-terms, we will define two closely related rewrite systems ***Reg*** and ***Reg*$^+$**. Rewrite sequences in ***Reg*** and in ***Reg*$^+$** deconstruct λ-terms by steps that typically decompose applications and λ-abstractions that occur just below the abstraction prefix. ***Reg*** and ***Reg*$^+$** differ with respect to the steps for removing vacuous prefix bindings they facilitate: while such bindings can always be removed by pertinent steps in ***Reg***, the system ***Reg*$^+$** only enables steps that remove vacuous bindings at the end of the abstraction prefix.

§ 1.2.11 (scope-delimiting strategies)**.** For each of these systems we consider a family of strategies that make deterministic choices concerning the application of the steps for removing vacuous prefix bindings: we call these strategies 'scope-delimiting strategies' for ***Reg***, and 'scope$^+$-delimiting strategies'[2] for ***Reg*$^+$**. Scope-delimiting strategies $\mathbb{S}$ for ***Reg*** and scope$^+$-delimiting strategies $\mathbb{S}^+$ for ***Reg*$^+$** induce rewrite relations $\to_{\mathbb{S}}$ and $\to_{\mathbb{S}^+}$, respectively. These families of rewrite strategies define respective notions of 'generated subterm', and they give rise to differently strong concepts of regularity: a λ-term M is called regular (strongly regular) if there is a rewrite strategy $\mathbb{S}$ for ***Reg*** (a rewrite strategy $\mathbb{S}^+$ for ***Reg*$^+$**) such that the set of from M $\to_{\mathbb{S}}$-reachable ($\to_{\mathbb{S}^+}$-reachable) generated subterms is finite.

[2] We usually say 'extended scope' when pronouncing scope$^+$, see § 1.4.5

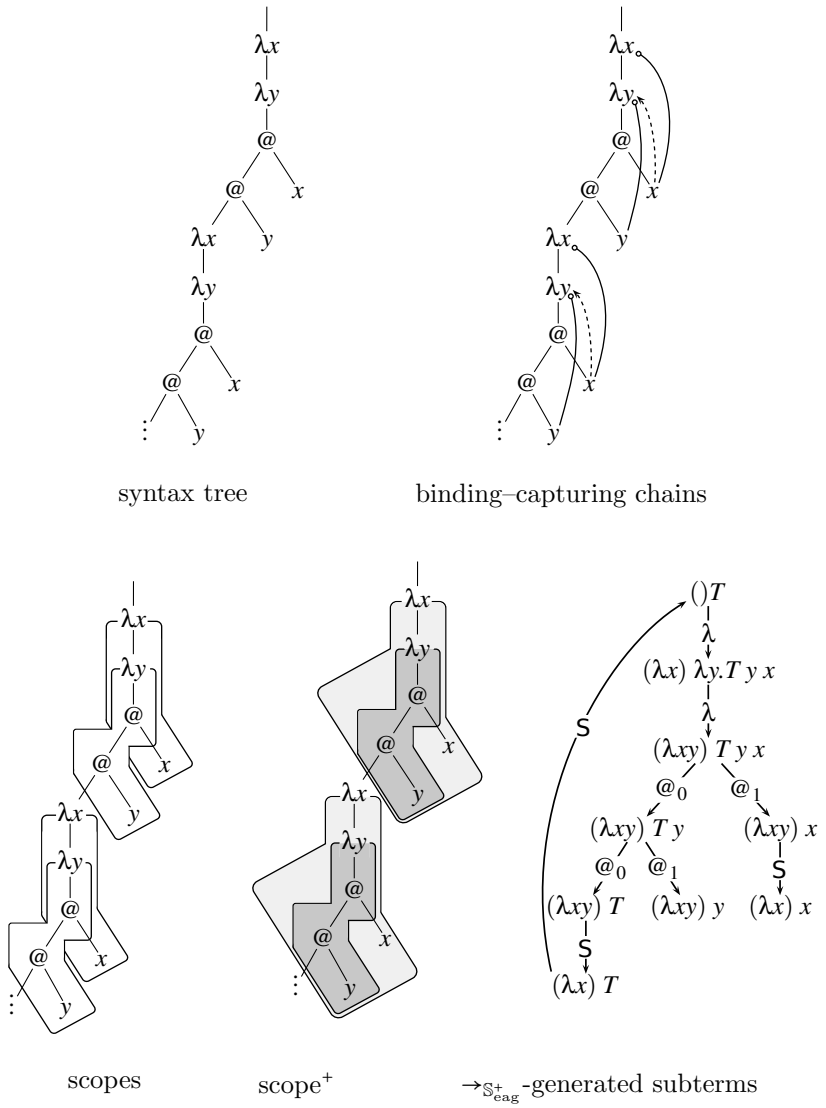

Figure 1.4. Various depictions of the strongly regular infinite λ-term that is expressed by the $\lambda_{\mathsf{letrec}}$-term $\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f$.

Example 1.2.12 (***Reg*$^+$-decomposition)**.** Before giving definitions of the decomposition CRSs (introduced in section 1.4) let us consider an example and decompose[3] the λ-term M of the form $M = \lambda xy.\, M\ y\ x$ from example 1.2.4. The set of generated subterms consists of these prefixed terms (in no particular order):
$()\ M \quad (\lambda x)\ \lambda y.\, M\ y\ x \quad (\lambda xy)\ M\ y\ x \quad (\lambda xy)\ M\ y \quad (\lambda xy)\ M \quad (\lambda x)\ M$
$(\lambda xy)\ y \quad (\lambda xy)\ x \quad (\lambda x)\ x.$
Here, $(\lambda xy)\ M\ y$ for instance is a witness for $M\ y$ being a (generated) subterm of M, with the variables x and y being bound 'somewhere above'. We could of course also have written $(\lambda xz)\ M\ z$, which is equivalent. So why is the subterm $(\lambda xy)\ x$ in the list along with $(\lambda x)\ x$, if x occurs at a position where both x and y have been bound? It is, because the decomposition rewrite systems include rules to remove variables from the prefix when no longer required (for leaving the scope). See fig. 1.4 for the reduction graph (Terminology 1.3.10) of M, which illustrates how the decomposition deconstructs M into the generated subterms above. Since $()\ M$ has only 9 different $\to_{\mathbb{S}^+_{\mathrm{eag}}}$-reducts, the term M is strongly regular.

§ 1.2.13 (result)**.** The generalisations of the concept of regularity to λ-terms suggest the question: do the expressibility results in [13] for regular first-order trees with respect to systems of recursion equations, rational expressions, or μ-terms also generalise in an appropriate way? We tackle only the case of strong regularity here, and obtain an expressibility result with respect to the $\boldsymbol{\lambda}_{\mathsf{letrec}}$, the λ-calculus with letrec. We show that an infinite unfolding is unique if it exists, and it can be obtained as the limit of an infinite rewrite sequence of unfolding steps. We prove that a λ-term is $\lambda_{\mathsf{letrec}}$-expressible if and only if it is strongly regular. This result confirms a conjecture[4] by Blom in [8, Section 1.2.4].

§ 1.2.14 (chapter overview)**.** section 1.3 introduces formalisms used in this chapter, predominantly concerning abstract rewriting systems and rewriting strategies. In section 1.4 we introduce rewriting systems for decomposing λ-terms into 'generated subterms', and we show some properties of these systems

[3] The term is decomposed with respect to the 'eager scope-delimiting' strategy $\mathbb{S}^+_{\mathrm{eag}}$ for ***Reg*$^+$**; see definition 1.4.34.

[4] Cf. the last sentence of [8, Section 1.2.4]: 'We conjecture that the set of regular lambda-trees is precisely the set of lambda-trees that can be obtained as the unwinding of terms with letrec'. Mind that 'regular lambda-trees' there correspond to strongly regular λ-terms in our sense, and that the notion of 'sub-tree' of a 'lambda-tree' corresponds to our notion of $\to_{\mathbb{S}^+}$-generated subterm with respect to a scope$^+$-delimiting strategy $\mathbb{S}^+$ for ***Reg*$^+$**.

in connection to so-called scope-delimiting strategies. Also in this section, we define regularity and strong regularity for λ-terms employing the concepts of generated subterms and scope-delimiting strategies. In section 1.5 we adapt the rewrite systems for decomposing λ-terms and the notions of scope-delimiting strategies to the λ-calculus with letrec. In section 1.6, we develop proof systems for the notions of regularity and strong regularity, for equality of strongly regular λ-terms, and for the property of a $\lambda_{\mathsf{letrec}}$-term to unfold to a λ-term. In section 1.7 we examine the binding structure of λ-terms (binding–capturing chains) and connect to the concepts introduced so far. In section 1.8 we establish the correspondence between strong regularity and $\lambda_{\mathsf{letrec}}$-expressibility for λ-terms. In section 1.9 we introduce 'λ-transition-graphs' of λ-terms and of $\lambda_{\mathsf{letrec}}$-terms as labelled transition graphs in which the edges carry one of the four different labels $@_0$, $@_1$, λ, and S. Section 1.10 summarises the above results.

1.3 Preliminaries

This section gathers known concepts vital to this chapter. Some notions concerning rewriting are recapitulated from [51], for others references are given. Some definitions of known concepts are simplified or tailored to our purposes.

Notation 1.3.1 ($|w|$, length of a word w)**.** For words w over some alphabet we denote the length of w by $|w|$.

We will use the following specific version of Kőnig's Lemma.

§ 1.3.2 (Kőnig's Lemma)**.** Let $G = \langle V, E\rangle$ be an undirected graph with set V of vertices and set E of edges. Suppose that G has infinitely many vertices (V is infinite), that it is connected (for all vertices $v, w \in V$ there exists a path in G from v to w) and that every vertex has finite degree (it is adjacent to only finitely many other vertices in G). Then for every vertex $v \in V$, G contains an infinitely long simple path from v, that is, a path starting at v without repetition of vertices.[5]

[5]This formulation corresponds to the following original formulation by Dénes Kőnig [38, page 80] "*Satz 3: Jeder unendliche zusammenhängende Graph G endlichen Grades besitzt einen einseitig unendlichen Weg, wobei der Anfangspunkt P_0 dieses Weges beliebig vorgeschrieben werden kann.*" in connection with the definition [38, page 10] "Eine unendliche Menge von Kanten P_iP_{i+1} ($i = 0, 1, \ldots$ *in inf.*), bzw. der durch sie gebildete Graph, heißt ein *einseitig unendlicher Weg*, falls für $i \neq j$ stets $P_i \neq P_j$ ist."

Rewriting Relations

Notation 1.3.3 ($R \cdot S$, composition of relations R and S)**.** For relations $R \subseteq A \times B$ and $S \subseteq B \times C$ we denote by $R \cdot S$ the *composition of R with S* defined by $R \cdot S := \{\langle x, z\rangle \mid (\exists y \in B)\langle x, y\rangle \in R \wedge \langle y, z\rangle \in S\} \subseteq A \times C$.

Notation 1.3.4 (R^*, reflexive transitive closure of a relation R)**.** For a relation $R \subseteq A \times B$ we denote by R^* the reflexive transitive closure of R under composition, for which it holds that $R^* := \bigcup_{i \in \mathbb{N}} R^i$ where $R^0 := \mathrm{id}_A := \{\langle x, x\rangle \mid x \in A\}$ and, for all $i \in \mathbb{N}$, $R^{i+1} := R \cdot R^i$.

Abstract Rewriting Systems

We use abstract rewriting systems (ARSs) to reason about CRSs (from which we derive the ARSs) and more specifically about the set of terms that a term can be reduced to in an arbitrary number of reductions. ARSs are essentially binary relations on sets. They are in a sense graph-like structures, and more tangible in comparison to CRSs and therefore easier to manipulate.

§ 1.3.5 (abstract rewriting systems)**.** An *abstract rewriting system (ARS)* is a quadruple $\langle O, \Phi, \mathrm{src}, \mathrm{tgt}\rangle$ consisting of a set O of *objects*, a set Φ of *steps*, and $\mathrm{src}, \mathrm{tgt} : \Phi \to O$, the *source* and *target* functions. For objects $o \in O$ we denote by $\Phi_{\mathrm{out}}(o)$ and by $\Phi_{\mathrm{in}}(o)$ the set of steps in Φ that depart (are outgoing steps) from o, and that arrive (are incoming steps) at o, respectively. We say that an ARS is *finite* if Φ is finite.

§ 1.3.6 (ARS induced by a CRS/iCRS)**.** Let $\mathcal{C}$ be a CRS/iCRS over signature Σ with rules R and let $\mathrm{Ter}(\Sigma)/\mathrm{Ter}^\infty(\Sigma)$ be the set of terms over Σ. We call the ARS $\mathcal{A} = \langle O, \Phi, \mathrm{src}, \mathrm{tgt}\rangle$ the *ARS induced by* $\mathcal{C}$ where:

$$O = \mathrm{Ter}(\Sigma) \quad / \quad O = \mathrm{Ter}^\infty(\Sigma)$$
$$\Phi = \{\langle s, \rho, t\rangle \mid s, t \in O, \rho \in R, s \to_\rho t\}$$
$$\mathrm{src} : \langle s, \rho, t\rangle \mapsto s$$
$$\mathrm{tgt} : \langle s, \rho, t\rangle \mapsto t$$

Notation 1.3.7 (induced ARS steps)**.** While the ARS formalism makes no qualitative distinction between different kinds of steps (Φ being an unstructured set in the definition of ARS), in § 1.3.6 we retain the original information from the CRS $\mathcal{C}$ as to which rule a step stems from (Φ being a set of triples). This allows us to do more fine-grained reasoning and we will continue to write

$o_1 \to_\rho o_2$ to indicate that there is a step $\phi \in \Phi$ with $\mathrm{src}(\phi) = o_1$ and $\mathrm{tgt}(\phi) = o_2$ which is due to the application of a CRS rule $\varrho^\rho_{\mathcal{C}}$, or even $o_1 \to_{\mathcal{A}.\rho} o_2$ to indicate which ARS we mean specifically. We sometimes also name the step specifically and write $\phi : o_1 \to_\rho o_2$ or $\phi : o_1 \to_{\mathcal{A}.\rho} o_2$.

§ 1.3.8 (sub-ARS)**.** Let $\mathcal{A}_1 = \langle O_1, \Phi_1, \mathrm{src}_1, \mathrm{tgt}_1\rangle$ and $\mathcal{A}_2 = \langle O_2, \Phi_2, \mathrm{src}_2, \mathrm{tgt}_2\rangle$ be ARSs. We say that $\mathcal{A}_1$ is a *sub-ARS* of $\mathcal{A}_2$ if $O_1 \subseteq O_2$, $\Phi_1 \subseteq \Phi_2$, and src_1, tgt_1 are the restrictions of src_2 and tgt_2, respectively, to Φ_1, which are required to be total functions. This implies that, for all $\phi \in \Phi_1$, it holds that $\mathrm{src}_2(\phi) = \mathrm{src}_1(\phi) \in O_1$, and $\mathrm{tgt}_2(\phi) = \mathrm{tgt}_1(\phi) \in O_1$.

§ 1.3.9 (generated sub-ARS)**.** For an object $o \in O$ of an ARS $\mathcal{A} = \langle O, \Phi, \mathrm{src}, \mathrm{tgt}\rangle$ we denote by $(o \twoheadrightarrow) := \langle O', \Phi', \mathrm{src}', \mathrm{tgt}'\rangle$ the *sub-ARS of* $\mathcal{A}$ *generated by* o, where O' comprises only the objects from O that are reachable from o by an arbitrary number of steps (or no steps) and with Φ', src', tgt' being the restrictions of Φ, src, tgt to the objects in O' and to steps between objects in O'. We write $(o \twoheadrightarrow_{\mathcal{A}})$ to explicitly refer to a specific ARS.

Terminology 1.3.10 (reduction graph)**.** We call $(o \twoheadrightarrow)$ the *reduction graph* of o, and $(o \twoheadrightarrow_{\mathcal{A}})$ the reduction graph of o with respect to $\mathcal{A}$.

ARS Strategies

Next we will give definitions for history-free and history-aware strategies for ARSs. The latter is based on the notion of ARS labelling which in turn is based on the notion of ARS bisimulations. Note that ARS bisimulation is solely introduced for the sake of defining ARS labellings and is not to be confused with bisimulation between LTSs or TRSs.

§ 1.3.11 (bisimulation between ARSs)**.** Let $\mathcal{A}_i = \langle O_i, \Phi_i, \mathrm{src}_i, \mathrm{tgt}_i\rangle$ for $i \in \{1, 2\}$ be ARSs. A relation $\mathbb{B} \subseteq (O_1 \times O_2) \cup (\Phi_1 \times \Phi_2)$, which relates objects with objects and steps with steps, is called an *ARS bisimulation* between $\mathcal{A}_1$ and $\mathcal{A}_2$ if the following holds:

- if $\mathbb{B}$ relates two objects, then $\mathbb{B}$ also relates their outgoing steps:

$$\begin{aligned}\forall \langle o_1, o_2\rangle \in O_1 \times O_2 \quad & o_1 \,\mathbb{B}\, o_2 \Rightarrow \\ & \forall \phi_1 \in \Phi_{1\mathrm{out}}(o_1)\ \exists \phi_2 \in \Phi_{2\mathrm{out}}(o_2) \quad \phi_1 \,\mathbb{B}\, \phi_2 \quad \wedge \\ & \forall \phi_2 \in \Phi_{2\mathrm{out}}(o_2)\ \exists \phi_1 \in \Phi_{1\mathrm{out}}(o_1) \quad \phi_2 \,\mathbb{B}\, \phi_1\end{aligned}$$

- if $\mathbb{B}$ relates two steps, then $\mathbb{B}$ also relates their sources and targets:

$$\forall \langle \phi_1, \phi_2 \rangle \in \Phi_1 \times \Phi_2 \quad \phi_1 \mathbb{B} \phi_2 \Rightarrow \mathrm{src}_1(\phi_1) \ \mathbb{B} \ \mathrm{src}_2(\phi_2) \wedge \mathrm{tgt}_1(\phi_1) \ \mathbb{B} \ \mathrm{tgt}_2(\phi_2)$$

In this work we need ARS-bisimulation only to define ARS-labellings:

§ 1.3.12 (labellings of ARSs)**.** Let $\mathcal{A} = \langle O, \Phi, \mathrm{src}, \mathrm{tgt} \rangle$ and $\mathcal{A}' = \langle O', \Phi', \mathrm{src}', \mathrm{tgt}' \rangle$ be ARSs.

(i) An ARS bisimulation $\mathbb{L}$ between $\mathcal{A}$ and $\mathcal{A}'$ is called a *labelling of* $\mathcal{A}$ *to* $\mathcal{A}'$, and $\mathcal{A}'$ *the* $\mathbb{L}$*-labelled version of* $\mathcal{A}$, if the converse $\mathbb{L}^\smile$ of $\mathbb{L}$ is a function $\mathbb{L}^\smile : O' \cup \Phi' \to O \cup \Phi$, and if additionally, for all $o' \in O'$ and $o \in O$ with $o \mathbb{L} o'$, the restriction $\mathbb{L}^\smile|_{\Phi'_{\mathrm{out}}(o')} \colon \Phi'_{\mathrm{out}}(o') \to \Phi_{\mathrm{out}}(o)$ of $\mathbb{L}^\smile$ to the steps departing from o' is bijective.

(ii) A *rewrite labelling* L of $\mathcal{A}$ to $\mathcal{A}'$ is a pair $\langle \mathbb{L}, l \rangle$ consisting of a labelling $\mathbb{L}$ of $\mathcal{A}$ to $\mathcal{A}'$ together with an *initial* labelling function l mapping objects of $\mathcal{A}$ to bisimilar objects of $\mathcal{A}'$.

§ 1.3.13 (history-free strategy)**.** A *history-free strategy* for an abstract rewriting system $\mathcal{A}$ is a sub-ARS of $\mathcal{A}$ that has the same objects, and the same normal forms as $\mathcal{A}$.

§ 1.3.14 (history-aware strategy)**.** A *history-aware strategy* for an abstract rewriting system $\mathcal{A}$ is a history-free strategy for the $\mathbb{L}$-labelled version of $\mathcal{A}$ with respect to, and together with, a rewrite labelling $\langle \mathbb{L}, l \rangle$ of $\mathcal{A}$.

§ 1.3.15 (history-aware and history-free strategies)**.** While strategies are simply defined as sub-ARSs of an ARS, here one may think of them as restrictions on the applicability of CRS rules. This is because in this work we consider ARSs that are induced by a CRS. A history-free strategy cannot restrict the applicability of a CRS rule to some term differently depending on the history (i.e. the previously applied rules) of that term. In history-aware strategy, on the other hands, terms are coupled with additional information that can be used to record the history of a term. Thus, how the applicability of CRS rules are restricted can differ depending on that record.

By a *strategy* for $\mathcal{A}$ we will mean a history-free strategy or a history-aware strategy for $\mathcal{A}$.

§ 1.3.16 (projecting history-aware strategies to history-free strategies)**.** Let $\mathbb{S}$ be a history-aware strategy for $\mathcal{A}$, and let $\mathcal{A}'$ be that $\mathbb{L}$-labelled version of $\mathcal{A}$ which $\mathbb{S}$ is a history-free strategy for. Then $\mathbb{S}$ projects to a history-free strategy $\check{\mathbb{S}}$ of $\mathcal{A}$. The projection is defined by $\mathbb{L}$, which induces a local bijective correspondence between outgoing steps of related sources of $\mathcal{A}$ and $\mathcal{A}'$. Mind that for deterministic $\mathbb{S}$, $\check{\mathbb{S}}$ may become non-deterministic. Furthermore, every rewrite sequence according to $\mathbb{S}$ in $\mathcal{A}'$ projects to a unique rewrite sequence in $\mathcal{A}$ (which is a rewrite sequence according to $\check{\mathbb{S}}$).

The last mentioned fact makes it possible to speak, for a given rewrite labelling, of rewrite sequences of a history-aware strategy for the objects of the original ARS.

Let $\mathbb{S}$ be a history-aware strategy for an ARS $\mathcal{A}$, and o an object of $\mathcal{A}$. Suppose that $\mathbb{S}$ is a sub-ARS of the $\mathbb{L}$-labelled version $\mathcal{A}'$ of $\mathcal{A}$ for some rewrite labelling $\langle \mathbb{L}, l \rangle$ of $\mathcal{A}$. Then by a *rewrite sequence of* $\mathbb{S}$ *on* o (in $\mathcal{A}$) we will mean the projection to $\mathcal{A}$ of a rewrite sequence of $\mathbb{S}$ (in $\mathcal{A}'$) on the result $l(o)$ of the initial labelling applied to o.

Labelled Transition Systems

§ 1.3.17 (labelled transition systems)**.** A *labelled transition system (LTS)* is a triple $\mathcal{L} = \langle S, L, T \rangle$ consisting of a set S of *states*, a set L of *labels*, and a set $T \subseteq S \times L \times S$ of L-labelled transitions. We write $s_1 \to_a s_2$ to denote labelled transitions, indicating that $\langle s_1, a, s_2 \rangle \in T$.

§ 1.3.18 (bisimulation between LTSs)**.** Let $\mathcal{L}_1 = \langle S_1, L, T_1 \rangle$ and $\mathcal{L}_2 = \langle S_2, L, T_2 \rangle$ be a LTSs over a common set of labels. A *bisimulation on* $\mathcal{L}$ is a binary relation $R \subseteq S_1 \times S_2$ that satisfies, for all $s_1 \in S_1$ and $s_2 \in S_2$:

(i) if $s_1 \; R \; s_2$ and $s_1 \to_a s_1'$, then there exists $s_2' \in S_2$ such that $s_2 \to_a s_2'$ and $s_1' \; R \; s_2'$;

(ii) if $s_1 \; R \; s_2$ and $s_2 \to_a s_2'$, then there exists $s_1' \in S_1$ such that $s_1 \to_a s_1'$ and $s_1' \; R \; s_2'$.

Two states $s_1 \in S_1$ and $s_2 \in S_2$ are *bisimilar*, denoted by $s_1 \leftrightarroweq s_2$, if there exists a bisimulation R such that $s_1 \; R \; s_2$.

Remark 1.3.19 (ARSs versus LTSs)**.** Note that labelled transition systems (LTSs) are essentially the same as (indexed) ARSs [51, 1.1]. Traditionally, research about LTSs is more concerned with the transitions and their labels

while research about ARSs is more concerned with set of objects and their reachability. This is reflected in a different definition of bisimulation for the two systems. Only on LTSs the bisimulation is sensitive to transition labels; ARSs do not have explicit labels on their steps. In this work we adhere to the tradition. We use ARSs for the definition of (strong) regularity which depends on whether the set of reachable terms is finite, and we use LTSs to define a graph representation for λ-terms on which we study properties revolving around (functional) bisimulation.

§ 1.3.20 (labelled transition graphs)**.** A *labelled transition graph (LTG)* G is a pointed LTS, that is, $G = \langle S, L, i, T\rangle$ where $\langle S, L, T\rangle$ is an LTS, and $i \in S$, which is called the *initial state*.

§ 1.3.21 (bisimulation between LTGs)**.** Two LTGs $G_1 = \langle S_1, L, i_1, \to_1\rangle$ and $G_1 = \langle S_2, L, i_2, \to_2\rangle$ are *bisimilar* if there is a bisimulation on the underlying LTSs that relates the initial states i_1 and i_2 with one another.

iCRS Terms

§ 1.3.22 (iCRS preterms, iCRS terms)**.** When speaking of 'infinite terms' for CRSs over some signature we draw on [51, 12.4] and [31] where metaterms of iCRSs are defined by means of metric completion. The metric is defined on α-equivalence classes of finite metaterms dependent on the minimal depth at which two finite preterms belonging to the equivalence classes have a 'conflict'. *iCRS terms* are defined as the objects formed by the metric completion process. They can be represented as equivalence classes of infinite preterms, which we call *iCRS preterms*, with respect to a notion of α-equivalence that again is based on the notion of 'conflict' (see also [51, Definition 12.4.1]). iCRS preterms are infinite ordered dyadic trees in which each node is either labelled by a variable name, and then the node does not have a successor, or by named abstractions λx (with some variable name x), and then the node has a single successor node, or by an application symbol, and then the node has a right and a left successor node.

Definition 1.3.23 (α-equivalence for iCRS preterms, Schroer-style proof system)**.** The notion of α-equivalence on iCRS preterms based on the absence of conflicts can be described by provability in the proof system $\mathbf{A}^{\infty}_{\mathcal{S}}(\Sigma)$ in fig. 1.5 which is an adaptation of a proof system due to Schroer [26]. The proof system $\mathbf{A}^{\infty}_{\mathcal{S}}(\Sigma)$ over signature Σ consists of the axioms and rules displayed in fig. 1.5 and contains a rule f for every $f \in \Sigma$. Provability in $\mathbf{A}^{\infty}_{\mathcal{S}}(\Sigma)$ of an equation

$$\frac{}{c = c}\ \text{const} \qquad \frac{s[x := c] = t[y := c]}{[x]s = [y]t}\ [\,]$$

$$\frac{s_1 = t_1 \quad \ldots\ldots \quad s_n = t_n}{f(s_1, \ldots, s_n) = f(t_1, \ldots, t_n)}\ f$$

Figure 1.5. Schroer-style proof system $\mathbf{A}^{\infty}_{\mathcal{S}}(\Sigma)$ for α-equivalence of iCRS preterms over signature Σ: for every $f \in \Sigma$ with arity n, $\mathbf{A}^{\infty}_{\mathcal{S}}(\Sigma)$ contains a rule f. In instances of the rule $[\,]$, the constant c is chosen fresh for s and t. Substitution which occurs in the assumption of $[\,]$ denotes substitution by variable replacement on iCRS preterms. It needs not to be capture-avoiding because of the freshness of the substituant.

between preterms is defined as the existence of a possibly infinite, *completed* derivation: for example, by ${}^{\infty}\vdash_{\mathbf{A}^{\infty}_{\mathcal{S}}} s = t$ we mean the existence of a possibly infinite proof tree $\mathcal{D}^{\infty}$ with conclusion $s = t$ such that maximal threads from the conclusion upwards either have length ω, or have finite length and end at a leaf that carries an axiom. (We will generally use the decorated turnstyle symbol ${}^{\infty}\vdash$ to indicate provability by a completed, possibly infinite derivation.)

However, closer to coinductive proof systems for λ-terms that we develop is the following different, but equivalent characterisation of α-equivalence for infinite iCRS preterms, a variant for iCRS terms of a proof system for α-equivalence between finite λ-terms due to Kahrs (see [26]).

Definition 1.3.24 (α-equivalence for iCRS preterms, Kahrs-style proof system)**.** The proof system $\mathbf{A}^{\infty}_{\mathcal{K}}(\Sigma)$ for α-equivalence on iCRS preterms over signature Σ consists of the axioms and the rules in fig. 1.6 with, for every $f \in \Sigma$, a rule f. Provability of an equation between preterms in $\mathbf{A}^{\infty}_{\mathcal{K}}(\Sigma)$ is defined, analogously as in fig. 1.5, as the existence of a possibly infinite, completed derivation.

This formulation of α-equivalence for λ-terms will be the key to our formulation of a 'coinduction principle' for λ-terms in theorem 1.9.12.

§ 1.3.25 (infinite rewrite relation, $\twoheadrightarrow$)**.** For an iCRS with rewrite relation $\to$, we denote by $\twoheadrightarrow$ the infinite rewrite relation induced by strongly convergent and continuous rewrite sequences of arbitrary (countable) ordinal length. Hereby strong convergence means that, at every limit ordinal, the depth of the rewrite

$$\frac{}{\{\vec{x}y\}\,y = \{\vec{z}u\}\,u}\ \mathsf{0}$$

$$\frac{\{\vec{x}\}\,s = \{\vec{z}\}\,t}{\{\vec{x}y\}\,s = \{\vec{z}w\}\,t}\ \mathsf{S}\quad \begin{array}{l}(\text{if } y \text{ does not occur in } s,\\ \text{and } w \text{ does not occur in } t)\end{array}$$

$$\frac{\{\vec{x}y\}\,s = \{\vec{z}u\}\,t}{\{\vec{x}\}\,[y]\,s = \{\vec{z}\}\,[u]\,t}\ [\,]$$

$$\frac{\{\vec{x}\}\,s_1 = \{\vec{y}\}\,t_1 \qquad \ldots\ldots \qquad \{\vec{x}\}\,s_n = \{\vec{y}\}\,t_n}{\{\vec{x}\}\,f(s_1,\ldots,s_n) = \{\vec{y}\}\,f(t_1,\ldots,t_n)}\ f$$

Figure 1.6. Kahrs-style proof system $\mathbf{A}^{\infty}_{\mathcal{K}}(\Sigma)$ for α-equivalence on iCRS preterms over signature Σ: for every $f \in \Sigma$ with arity n, $\mathbf{A}^{\infty}_{\mathcal{K}}(\Sigma)$ contains a rule f.

activity in the rewrite sequence's terms tends to infinity. Continuity means that the terms of the rewrite sequence converge, in the metric space of infinite terms, at every limit ordinal. By $\to^{\omega}$ we will denote the rewrite relation induced by strongly continuous $\to$-rewrite-sequences of length ω.

1.4 Regular and strongly regular λ-terms

§ 1.4.1 (regularity)**.** For infinite first-order trees the concept of regularity is well-known and well-studied [13]. Regularity of a labelled tree[6] is defined as the existence of only finitely many subtrees and implies the existence of a finite graph that unfolds to that tree. In this section we generalise the notion of regularity to trees with a binding mechanism, i.e. to λ-terms specifically. We give a definition for regularity which corresponds to regularity of a term when regarded as a first-order tree, and for strong regularity, which will be shown in the following sections to coincide with $\lambda_{\mathsf{letrec}}$-expressibility.

[6]By a 'labelled tree' we here mean a finite or infinite tree whose nodes are labelled by function symbols with a fixed arity. The arity determines the number of (ordered) successor nodes each node has.

§ 1.4.2 (decomposition CRSs ***Reg*** and ***Reg***$^{+}$)**.** We define regularity and strong regularity in terms of rewriting systems that will be called ***Reg*** and ***Reg***$^{+}$. Rewrite sequences in these systems *inspect* a given term coinductively in the sense that a rewrite sequence corresponds to a decomposition of the term along one of its paths from the root. Both ***Reg*** and ***Reg***$^{+}$ extend a kernel system ***Reg***$^{-}$ comprising three rewrite rules which denote whether the position just passed in the tree is an abstraction or an application and in the second case whether the application is being followed to the left or to the right.

§ 1.4.3 (prefixed terms)**.** The rewriting systems are defined on λ-terms enriched by what we call an *abstraction prefix*, by which the terms can be kept closed during the deconstruction. This is crucial for the definition of the rewriting system as a CRS. While intuitively it is clear that the λ-term $M\ N$ is composed of the subterms M and N, abstractions are more problematic. In a first-order setting one could say that $\lambda x.\, M$ contains M as a subterm, but if x occurs freely in M then M would be an open term. That means that the scrutinisation of an abstraction would be able to go from a closed term to an open term, which would run counter to the interpretation of a higher-order term as an α-equivalence class. In the definition of the rewriting systems below this issue is resolved as follows. When inspecting an abstraction, the binder is eliminated but moved from the scrutinised subterm into the prefix. That guarantees that the term as a whole remains closed.

Example 1.4.4. In $\lambda x.\, \lambda y.\, x\ x\ y$ for instance the path from the root to the second occurrence of x then corresponds to the rewrite sequence: $()\ \lambda x.\, \lambda y.\, x\ x\ y \to_{\lambda} (\lambda x)\ \lambda y.\, x\ x\ y \to_{\lambda} (\lambda xy)\ x\ x\ y \to_{@_0} (\lambda xy)\ x\ x \to_{@_1} (\lambda xy)\ x$.

§ 1.4.5 (scope and scope^{+})**.** The ***Reg*** and the ***Reg***$^{+}$ system extend ***Reg***$^{-}$ by a *scope-delimiting* rule, which signifies that the scope of an abstraction has ended, with both systems being based on different notions of scope. ***Reg*** relies on what we simply call *scope* of an abstraction: the range from the abstraction up to the positions under which the bound variable does not occur anymore. We base ***Reg***$^{+}$ on a different notion of scope, called scope^{+} (pronounced 'extended scope'), which is strictly nested. The scopes^{+} of an abstraction extends its scope by encompassing all scopes^{+} that are opened within its range. As a consequence, scopes^{+} do no overlap partially (see fig. 1.7) but are strictly nested. In a sense scope^{+} is the transitive closure of scope; this becomes obvious in definition 1.7.9 where a precise definition of scope and scope^{+} is given. In [5] scopes^{+} are called 'skeletons'.

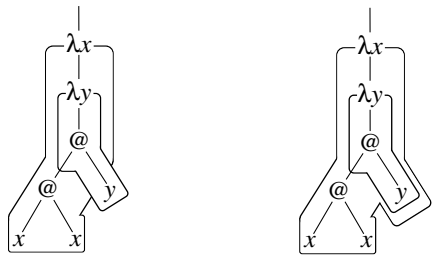

Figure 1.7. The difference between scope and scope$^+$

§ 1.4.6 (scope$^+$-delimiters identify abstractions)**.** When every scope$^+$ is closed by the scope$^+$-delimiting rule then the sequence of rewrite steps alone (i.e. without the terms themselves) unambiguously determines which abstraction a variable occurrence belongs to.

Example 1.4.7. The rewrite sequence from above would then have one additional scope$^+$-delimiting step asserting that the variable at the end of the path is indeed x and not y: $()\ \lambda x.\,\lambda y.\,x\ x\ y\ \to_\lambda\ (\lambda x)\ \lambda y.\,x\ x\ y\ \to_\lambda\ (\lambda xy)\ x\ x\ y \to_{@_0} (\lambda xy)\ x\ x \to_{@_1} (\lambda xy)\ x \to_{\mathsf{S}} (\lambda x)\ x$

§ 1.4.8 (abstraction prefix and scope$^+$)**.** The abstraction prefix not only keeps the term closed but also denotes which scope$^+$ is still open, which provides the information to decide applicability of the scope$^+$-delimiting rule.

Example 1.4.9. The last step closes the scope$^+$ of y, therefore that variable is removed from the prefix. The rewrite sequence for the path to the occurrence of y does not include a scope$^+$-delimiting step: $()\ \lambda x.\,\lambda y.\,x\ x\ y \to_\lambda (\lambda x)\ \lambda y.\,x\ x\ y \to_\lambda (\lambda xy)\ x\ x\ y \to_{@_1} (\lambda xy)\ y$

§ 1.4.10 (nameless representation for λ-terms)**.** Ultimately, the ***Reg*$^+$** rewriting system defines nameless representations for λ-terms related to the de-Bruijn notation for λ-terms. Considering the de-Bruijn representation of the above term $\lambda.\,\lambda.\,\mathsf{S}(0)\ \mathsf{S}(0)\ 0$ we find that the position of the $\to_{\mathsf{S}}$-steps indeed coincides with the position of the S markers. However, the rewrite system ***Reg*$^+$** permits more flexibility for the placement of $\to_{\mathsf{S}}$-steps.

Example 1.4.11. For example the path from example 1.4.7 to the second occurrence of x can also be witnessed by the following rewrite sequence: $()\ \ \lambda x.\,\lambda y.\,x\ x\ y\ \ \to_\lambda\ \ (\lambda x)\ \ \lambda y.\,x\ x\ y\ \ \to_\lambda\ \ (\lambda xy)\ \ x\ x\ y\ \ \to_{@_0}\ \ (\lambda xy)\ \ x\ x\ \ \to_{\mathsf{S}}$

$(\lambda x)\ x\ x \to_{@_1} (\lambda x)\ x$. Here the scope$^+$ of y is closed earlier. This would correspond to $\lambda.\lambda.\mathsf{S}(0\ 0)\ 0$ in de-Bruijn notation, or more precisely, in a variant of the de-Bruijn notation which permits the scope/scope$^+$-delimiter S to occur anywhere between a variable occurrence and its binding abstraction.

§ 1.4.12 (scope delimiters in the literature)**.** This variation of the de-Bruijn notation [10], in which S-symbols can be used anywhere in the term signify the end of a scope, is due to Paterson [7]. The idea is also used in [47] and related to an even more flexible end-of-scope symbol ↼ [26].

Definition 1.4.13 (CRS terms with abstraction prefixes)**.** The CRS signature for $\boldsymbol{(\lambda)}$, the λ-calculus with abstraction prefixes, extends the CRS signature Σ^λ for $\boldsymbol{\lambda}$ (see definition 0.5.4) and consists of the set $\Sigma^{(\lambda)} = \Sigma^\lambda \cup \{\mathsf{pre}_n \mid n \in \mathbb{N}\}$ of function symbols, where for $n \in \mathbb{N}$ the function symbols pre_n for *prefix* λ-*abstractions* of length n are unary (have arity one). CRS terms with leading prefixes $\mathsf{pre}_n([x_1]\dots[x_n]M)$ will informally be denoted by $(\lambda x_1 \dots x_n)\ M$, abbreviated as $(\lambda \vec{x})\ M$.

Definition 1.4.14 (the CRSs $\boldsymbol{Reg^-}$, $\boldsymbol{Reg}$, $\boldsymbol{Reg^+}$ for decomposing λ-terms)**.** Consider the following CRS rule schemes over $\Sigma^{(\lambda)}$:

$$
\begin{aligned}
\varrho_\Delta^{@_0} &: \mathsf{pre}_n([x_1 \dots x_n]\mathsf{app}(Z_0(\vec{x}), Z_1(\vec{x}))) \to \mathsf{pre}_n([x_1 \dots x_n]Z_0(\vec{x})) \\
\varrho_\Delta^{@_1} &: \mathsf{pre}_n([x_1 \dots x_n]\mathsf{app}(Z_0(\vec{x}), Z_1(\vec{x}))) \to \mathsf{pre}_n([x_1 \dots x_n]Z_1(\vec{x})) \\
\varrho_\Delta^{\lambda} &: \mathsf{pre}_n([x_1 \dots x_n]\mathsf{abs}([x_{n+1}]Z(\vec{x}))) \to \mathsf{pre}_{n+1}([x_1 \dots x_{n+1}]Z(\vec{x})) \\
\varrho_\Delta^{\mathsf{S}} &: \mathsf{pre}_{n+1}([x_1 \dots x_{n+1}]Z(x_1, \dots, x_n)) \to \mathsf{pre}_n([x_1 \dots x_n]Z(x_1, \dots, x_n)) \\
\varrho_\Delta^{\mathrm{del}} &: \mathsf{pre}_{n+1}([x_1 \dots x_{n+1}]Z(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{n+1})) \to \\
&\qquad \mathsf{pre}_n([x_1 \dots x_{i-1}x_{i+1} \dots x_n]Z(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{n+1}))
\end{aligned}
$$

By $\boldsymbol{Reg^-}$ we denote the CRS with rules $\varrho_\Delta^{@_0}$, $\varrho_\Delta^{@_1}$, and ϱ_Δ^{λ}. By $\boldsymbol{Reg}$ (by $\boldsymbol{Reg^+}$) we denote the CRS consisting of all of the above rules *except* the rule $\varrho_\Delta^{\mathsf{S}}$ (*except* the rule $\varrho_\Delta^{\mathrm{del}}$). The rewrite relations of $\boldsymbol{Reg^-}$, $\boldsymbol{Reg}$, and $\boldsymbol{Reg^+}$ are denoted by $\to_{\mathrm{reg}^-}$, $\to_{\mathrm{reg}}$ and $\to_{\mathrm{reg}^+}$, respectively. And by $\to_{@_0}$, $\to_{@_1}$, $\to_\lambda$, $\to_{\mathsf{S}}$, $\to_{\mathrm{del}}$, we respectively denote the rewrite relations induced by each of the single rules $\varrho_\Delta^{@_0}$, $\varrho_\Delta^{@_1}$, ϱ_Δ^{λ}, $\varrho_\Delta^{\mathsf{S}}$, and $\varrho_\Delta^{\mathrm{del}}$.

§ 1.4.15 ($\boldsymbol{Reg^-}$, $\boldsymbol{Reg}$, and $\boldsymbol{Reg^+}$ in informal notation)**.** For better readability we will from now on rely on the informal notation corresponding to definition 1.4.14

which is as follows:

$$\begin{aligned}
\varrho_{\Delta}^{@_0} &: (\lambda x_1 \dots x_n)\ M_0\ M_1 \to (\lambda x_1 \dots x_n)\ M_0 \\
\varrho_{\Delta}^{@_1} &: (\lambda x_1 \dots x_n)\ M_0\ M_1 \to (\lambda x_1 \dots x_n)\ M_1 \\
\varrho_{\Delta}^{\lambda} &: (\lambda x_1 \dots x_n)\ \lambda x_{n+1}.\, M_0 \to (\lambda x_1 \dots x_{n+1})\ M_0 \\
\varrho_{\Delta}^{\mathsf{S}} &: (\lambda x_1 \dots x_{n+1})\ M_0 \to (\lambda x_1 \dots x_n)\ M_0 \\
&\qquad \text{(if the binding } \lambda x_{n+1} \text{ is vacuous)} \\
\varrho_{\Delta}^{\mathrm{del}} &: (\lambda x_1 \dots x_{n+1})\ M_0 \to (\lambda x_1 \dots x_{i-1} x_{i+1} \dots x_{n+1})\ M_0 \\
&\qquad \text{(if the binding } \lambda x_i \text{ is vacuous)}
\end{aligned}$$

Remark 1.4.16. Be reminded that as mentioned in Notation 0.6.4 we will at times omit the Δ and write ϱ^{ρ} instead of ϱ_{Δ}^{ρ}.

§ 1.4.17 (***Reg***$^{+}$ defines nameless representations)**.** Considering the reduction graphs from fig. 1.8 without labels on their nodes we see that only from the ***Reg***$^{+}$ graph the original term can be reconstructed unambiguously. For example, the path to the rightmost occurrence of x has the rewrite sequence in ***Reg***

$$\to_{\lambda} \cdot \to_{\lambda} \cdot \to_{@_0} \cdot \to_{\mathrm{del}} \cdot \to_{@_1}$$

which witnesses an occurrence of y in place of x at the same position. This ambiguity plays a role for the definition of λ-transition-graphs in section 1.9, and is discussed in that context in § 1.9.14.

§ 1.4.18 ($\to_{\mathrm{reg}^+} \subseteq \to_{\mathrm{reg}} \subseteq \to_{\mathrm{reg}^-}$)**.** Note that $\to_{\mathrm{reg}}$ is contained in $\to_{\mathrm{reg}^-}$, and since the rule ϱ^{del} generalises the rule ϱ^{S}, $\to_{\mathrm{reg}^+}$ is contained in $\to_{\mathrm{reg}}$.

We will specifically need the following statement later on:

Proposition 1.4.19. Every rewrite sequence in ***Reg***$^{+}$ corresponds directly to a rewrite sequence in ***Reg*** by exchanging $\to_{\mathsf{S}}$-steps with $\to_{\mathrm{del}}$-steps.

§ 1.4.20 (abstract prefix symbol)**.** We are only interested in terms which have a single occurrence of the abstraction prefix symbol at the outermost position. Note that the rules in ***Reg*** and ***Reg***$^{+}$ guarantee that every reduct of a term of the form $(\lambda \vec{x})\ M$ is again a term of this form. Therefore we define:

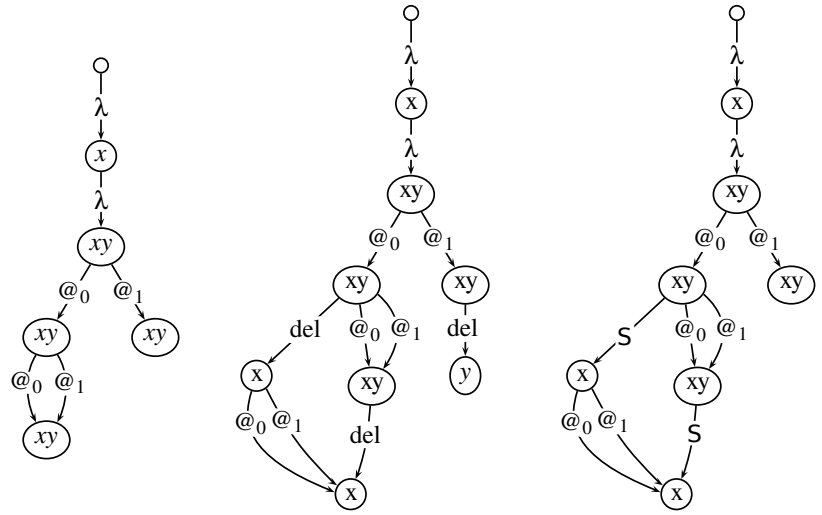

Figure 1.8. The reduction graphs of $M = () \; \lambda x.\, \lambda y.\, x\; x\; y$ with respect to $\boldsymbol{Reg^-}/\boldsymbol{Reg}/\boldsymbol{Reg^+}$. Note, that in the vertices not the entire $(\boldsymbol{\lambda})$-terms are displayed but only the respective prefixes.

Definition 1.4.21 (prefixed λ-terms)**.** By $\mathrm{Ter}^\infty((\boldsymbol{\lambda}))$ we denote the subset of closed iCRS terms over $\Sigma^{(\lambda)}$ with the restrictions:

- Every term $M \in \mathrm{Ter}^\infty((\boldsymbol{\lambda}))$ has a prefix at its root and nowhere else: M is of the form $\mathsf{pre}_n([x_1]\dots[x_n]M)$ and M has no occurrences of function symbols pre_i for any $i \in \mathbb{N}$.
- Otherwise a CRS abstraction can only occur directly beneath an abs-symbol.

$\mathrm{Ter}^\infty((\boldsymbol{\lambda}))$ is more formally specified in definition 1.6.3.

Proposition 1.4.22. $\mathrm{Ter}^\infty((\boldsymbol{\lambda}))$ is closed under $\to_{\mathrm{reg}^-}$, $\to_{\mathrm{reg}}$, and $\to_{\mathrm{reg}^+}$.

Definition 1.4.23 (the ARSs Reg^-, Reg, Reg^+)**.** We denote by Reg^-, Reg and Reg^+ the ARSs induced by the iCRSs derived from $\boldsymbol{Reg^-}$, $\boldsymbol{Reg}$, $\boldsymbol{Reg^+}$, restricted to terms in $\mathrm{Ter}^\infty((\boldsymbol{\lambda}))$.

Proposition 1.4.24. The restrictions of the rewrite relations of Reg^-, Reg and Reg^+ to $\mathrm{Ter}^\infty(\mathbf{(\lambda)})$, the set of objects of Reg^-, Reg, and Reg^+, have the following properties:

(i) $\to_{\mathrm{del}}$ is confluent, and terminating.

(ii) $\to_{\mathrm{del}}$ one-step commutes with $\to_\lambda$, $\to_{@_0}$, $\to_{@_1}$, $\to_{\mathsf{S}}$:

$$\leftarrow_{\mathrm{del}} \cdot \to_\rho \subseteq \to_\rho \cdot \leftarrow_{\mathrm{del}} \quad \forall \rho \in \{\lambda, @_0, @_1, \mathsf{S}\}$$

(iii) $\to_{\mathrm{del}}$ can be postponed:

$$\to_{\mathrm{del}} \cdot \to_\rho \subseteq \to_\rho \cdot \to_{\mathrm{del}} \quad \forall \rho \in \{\lambda, @_0, @_1, \mathsf{S}\}$$

(iv) $\to_{\mathsf{S}} \subseteq \to_{\mathrm{del}}$ and thus $\to_{\mathrm{reg}^+} \subseteq \to_{\mathrm{reg}}$. Furthermore $\to_{\mathrm{reg}^-} \subset \to_{\mathrm{reg}^+} \subset \to_{\mathrm{reg}}$.

(v) $\to_{\mathsf{S}}$ is deterministic, hence confluent, and terminating.

(vi) $\to_{\mathsf{S}}$ one-step commutes with $\to_{@_0}$, and $\to_{@_1}$:

$$\leftarrow_{\mathsf{S}} \cdot \to_{@_i} \subseteq \to_{@_i} \cdot \leftarrow_{\mathsf{S}} \quad \forall i \in \{0, 1\}$$

(vii) $(\lambda x)\ x$ is the sole term in $\to_{\mathrm{reg}}$-normal-form. $\to_{\mathrm{reg}^+}$-normal-forms are of the form $(\lambda x_1 \ldots x_n)\ x_n$.

(viii) $\to_{\mathrm{reg}}$ and $\to_{\mathrm{reg}^+}$ are finitely branching, and – on finite terms – terminating.

Proof. Most properties, including those concerning commutation of steps, are easy to verify by analysing the behaviour of the rewrite rules in Reg on terms of $\mathrm{Ter}^\infty(\mathbf{(\lambda)})$.

For $\to_{\mathrm{reg}^+} \subset \to_{\mathrm{reg}}$ in (iv) note that, for example, $(\lambda xy)\ y \to_{\mathrm{reg}} (\lambda y)\ y$ by a $\to_{\mathrm{del}}$-step, but that $(\lambda xy)\ y$ is a $\to_{\mathsf{S}}$-normal-form, and hence also a $\to_{\mathrm{reg}^+}$-normal-form.

Concerning (viii) we first argue for finite branchingness of $\to_{\mathrm{reg}}$ and $\to_{\mathrm{reg}^+}$ on $\mathrm{Ter}^\infty(\boldsymbol{\lambda})$: this property follows from the fact that, on a term $(\lambda \vec{x})\ M$ with just one abstraction in its prefix, of the constituent rewrite relations $\to_{@_0}$, $\to_{@_1}$, $\to_\lambda$, $\to_{\mathsf{S}}$, $\to_{\mathrm{del}}$ of $\to_{\mathrm{reg}}$ and $\to_{\mathrm{reg}^+}$ only $\to_{\mathrm{del}}$ can have branching degree greater than one, which in this case then also is bounded by the length $|\vec{x}|$ of the abstraction prefix. For termination of $\to_{\mathrm{reg}}$ and $\to_{\mathrm{reg}^+}$ on finite terms with just a leading abstraction prefix we can restrict to $\to_{\mathrm{reg}}$, due to (iv), and argue as follows: on finite terms in $\mathrm{Ter}^\infty(\mathbf{(\lambda)})$, in every $\to_{\mathrm{reg}}$-rewrite-step either the size

of the body of the term decreases strictly, or the size of the body stays the same, but the length of the prefix decreases by one. Hence in every rewrite step the measure $\langle \text{body size}, \text{prefix length} \rangle$ on terms decreases strictly in the (well-founded) lexicographic ordering on $\mathbb{N} \times \mathbb{N}$. □

Corollary 1.4.25 ($\to_{\mathrm{del}}$ and $\to_{\mathsf{S}}$ are normalising)**.** Note that as a consequence of proposition 1.4.24 (i) and (v), the rewrite relations $\to_{\mathrm{del}}$ and $\to_{\mathsf{S}}$ are normalising on $\mathrm{Ter}^{\infty}(\boldsymbol{\lambda})$.

Proposition 1.4.26. The following statements hold:

(i) Let $(\lambda\vec{x})\ M$ a term in *Reg* with $|\vec{x}| = n \in \mathbb{N}$. Then the number of terms $(\lambda\vec{y})\ N$ in *Reg* with $(\lambda\vec{y})\ N \twoheadrightarrow_{\mathrm{del}} (\lambda\vec{x})\ M$ and $|\vec{y}| = n + k \in \mathbb{N}$ is $\binom{n+k}{n}$.

(ii) Let T be a finite set of terms in *Reg*, and $k \in \mathbb{N}$. Then also the set of terms in *Reg* that are the form $(\lambda\vec{y})\ N$ with $|\vec{y}| \leq k$ and that have a $\twoheadrightarrow_{\mathrm{del}}$-reduct in T is finite.

Proof. From $(\lambda\vec{y})\ M(y) = (\lambda y_1 \dots y_{n+k})\ N(y_1, \dots, y_{n+k}) \twoheadrightarrow_{\mathrm{del}} (\lambda x_1 \dots x_n)\ M(x_1, \dots, x_n) = (\lambda\vec{x})\ M(\vec{x})$ it follows that there are $i_1, \dots, i_n \in \{1, \dots, n+k\}$ with $i_1 < i_2 < \dots < i_n$ such that the term $(\lambda\vec{y})\ M(y)$ is actually of the form $(\lambda y_1 \dots y_{n+k})\ N(y_{i_1}, \dots, y_{i_n})$ and furthermore $(\lambda y_{i_1} \dots y_{i_n})\ N(y_{i_1}, \dots, y_{i_n}) = (\lambda x_1 \dots x_n)\ M(x_1, \dots, x_n)$. Hence the number of terms $(\lambda\vec{y})\ M(y)$ with $\twoheadrightarrow_{\mathrm{del}}$-reduct $(\lambda\vec{x})\ M(\vec{x})$ is equal to the number of choices $i_1, \dots, i_n \in \{1, \dots, n+k\}$ such that $i_1 < i_2 < \dots < i_n$. This establishes statement (i). Statement (ii) is an easy consequence. □

Lemma 1.4.27. On $\mathrm{Ter}^{\infty}(\boldsymbol{(\lambda)})$, the rewrite relations $\to_{\mathrm{reg}}$ and $\to_{\mathrm{reg}^+}$ have the following further properties with respect to $\to_{\mathrm{del}}$, $\to_{\mathsf{S}}$, $\to^{!}_{\mathrm{del}}$, and $\to^{!}_{\mathsf{S}}$:

$$
\begin{aligned}
\twoheadleftarrow_{\mathrm{del}} \cdot \to_{\mathrm{reg}} &\subseteq (\to^{!}_{\mathrm{del}} \cdot \to^{=}_{\mathrm{reg}^-}) \cdot \twoheadleftarrow_{\mathrm{del}} \\
\twoheadleftarrow_{\mathrm{del}} \cdot \twoheadrightarrow_{\mathrm{reg}} &\subseteq (\to^{!}_{\mathrm{del}} \cdot \to^{=}_{\mathrm{reg}^-})^{*} \cdot \twoheadleftarrow_{\mathrm{del}} \\
\\
\twoheadleftarrow_{\mathrm{del}} \cdot \to_{\mathrm{reg}^+} &\subseteq (\to^{!}_{\mathsf{S}} \cdot \to^{=}_{\mathrm{reg}^-}) \cdot \twoheadleftarrow_{\mathrm{del}} \\
\twoheadleftarrow_{\mathrm{del}} \cdot \twoheadrightarrow_{\mathrm{reg}^+} &\subseteq (\to^{!}_{\mathsf{S}} \cdot \to^{=}_{\mathrm{reg}^-})^{*} \cdot \twoheadleftarrow_{\mathrm{del}}
\end{aligned}
$$

Proof. These commutation properties, which can be viewed as projection properties, can be shown by arguments with diagrams using the commutation properties in proposition 1.4.24. □

Remark 1.4.28. The commutation properties in lemma 1.4.27 can be refined to state that $\to_\lambda$-steps project to $\to_\lambda$-steps, and $\to_{@_0}$-steps and $\to_{@_1}$-steps project to $\to_{@_0}$-steps and $\to_{@_1}$-steps, accordingly.

As an immediate consequence of lemma 1.4.27 we obtain the following lemma, which formulates a connection via projection between rewrite sequences in *Reg* (in Reg^+) and $\to_{\mathrm{del}}$-*eager* ($\to_{\mathsf{S}}$-*eager*) rewrite sequences in *Reg* (in Reg^+) that do not contain $\to_\lambda$-steps, $\to_{@_0}$-steps, or $\to_{@_1}$-steps on terms that allow $\to_{\mathrm{del}}$-steps ($\to_{\mathsf{S}}$-steps).

Lemma 1.4.29. The following statements hold:

(i) Every (finite or infinite) rewrite sequence in *Reg* of the form

$$\tau : (\lambda\vec{x}_0)\ M_0 \to_{\mathrm{reg}} (\lambda\vec{x}_1)\ M_1 \to_{\mathrm{reg}} \ldots \to_{\mathrm{reg}} (\lambda\vec{x}_k)\ M_k \to_{\mathrm{reg}} \ldots$$

projects over a rewrite sequence $\pi : (\lambda\vec{x}_0)\ M_0 \twoheadrightarrow_{\mathrm{del}} (\lambda\vec{x}'_0)\ M_0$ to a $\to_{\mathrm{del}}$-eager rewrite sequence in *Reg* of the form

$$\begin{aligned}\check{\tau} : (\lambda\vec{x}'_0)\ M_0 \to^{!}_{\mathrm{del}} \cdot \to^{=}_{\mathrm{reg}^-} (\lambda\vec{x}'_1)\ M_1 \to^{!}_{\mathrm{del}} \cdot \to^{=}_{\mathrm{reg}^-} \ \ldots \\ \ldots \to^{!}_{\mathrm{del}} \cdot \to^{=}_{\mathrm{reg}^-} (\lambda\vec{x}'_k)\ M_k \to^{!}_{\mathrm{del}} \cdot \to^{=}_{\mathrm{reg}^-} \ \ldots\end{aligned}$$

in the sense that $(\lambda\vec{x}_i)\ M_i \twoheadrightarrow_{\mathrm{del}} (\lambda\vec{x}'_i)\ M_i$ for all $i \in \mathbb{N}$ less or equal to the length of τ.

(ii) Every (finite or infinite) rewrite sequence in Reg^+ of the form

$$\tau : (\lambda\vec{y}_0)\ N_0 \to_{\mathrm{reg}^+} (\lambda\vec{y}_1)\ N_1 \to_{\mathrm{reg}^+} \ldots \to_{\mathrm{reg}^+} (\lambda\vec{y}_k)\ N_k \to_{\mathrm{reg}^+} \ldots$$

projects over a rewrite sequence $\pi : (\lambda\vec{y}_0)\ N_0 \twoheadrightarrow_{\mathrm{del}} (\lambda\vec{y}'_0)\ N_0$ to a $\to_{\mathsf{S}}$-eager rewrite sequence in Reg^+:

$$\begin{aligned}\check{\tau} : (\lambda\vec{y}'_0)\ N_0 \to^{!}_{\mathsf{S}} \cdot \to^{=}_{\mathrm{reg}^-} (\lambda\vec{y}'_1)\ N_1 \to^{!}_{\mathsf{S}} \cdot \to^{=}_{\mathrm{reg}^-} \ \ldots \\ \ldots \to^{!}_{\mathsf{S}} \cdot \to^{=}_{\mathrm{reg}^-} (\lambda\vec{y}'_k)\ N_k \to^{!}_{\mathsf{S}} \cdot \to^{=}_{\mathrm{reg}^-} \ \ldots\end{aligned}$$

in the sense that $(\lambda\vec{y}_i)\ N_i \twoheadrightarrow_{\mathsf{S}} (\lambda\vec{y}'_i)\ N_i$ for all $i \in \mathbb{N}$ less or equal to the length of τ.

§ 1.4.30 (non-determinism of $\to_{\mathrm{reg}}$ and $\to_{\mathrm{reg}^+}$)**.** On terms in $\mathrm{Ter}^\infty((\boldsymbol{\lambda}))$, which have just one prefix at the top of the term, there are two different causes for non-determinism of the rewrite relations in **Reg** and $\mathbf{Reg^+}$: First, since the

left-hand sides of the rules $\varrho^{@_0}$ and $\varrho^{@_1}$ coincide, these rules enable different steps on the same term, producing the left- and respectively the right subterm of the application immediately below the prefix. Second, the rules ϱ^{del} and ϱ^{S} can be applicable in situations where also one of the rules $\varrho^{@_0}$, $\varrho^{@_1}$, or ϱ^{λ} is applicable (see for instance fig. 1.8, in the middle). Whereas the first kind of non-determinism is due to the 'observer' having to observe the two different subterms of an application in a λ-term, the second is due to a freedom of the observer as to when to attest the end of a scope/scope$^+$ in the analysed λ-term.

§ 1.4.31 (outlook: λ-transition-graphs)**.** In the definition below we define strategies for *Reg* and *Reg*$^+$ that only allow the former source of non-determinism while forbidding the second kind. The intention is that the reduction graph of a term with respect to such a strategy corresponds to the term's syntax tree. These reduction graphs can be seen as a nameless graph representation for λ-terms. We will introduce and study such representations (called λ-transition-graphs) in section 1.9 and a further adaptation (λ-term-graphs) in chapter 2.

Definition 1.4.32 (scope/scope$^+$-delimiting strategy)**.** We call a strategy $\mathbb{S}$ for *Reg* (for *Reg*$^+$) a *scope-delimiting strategy* (*scope*$^+$*-delimiting strategy*) if the source of a step is non-deterministic (that is, it is the source of more than one step) if and only if it is the source of precisely a $\rightarrow_{@_0}$-step and a $\rightarrow_{@_1}$-step.

§ 1.4.33 (alternative formulation of definition 1.4.32)**.** We can define a bit more verbosely: a strategy $\mathbb{S}$ for *Reg* (for *Reg*$^+$) is called a scope-delimiting (scope$^+$-delimiting) strategy if

- every source of a step is one of three kinds: the source of a $\rightarrow_\lambda$-step, the source of a $\rightarrow_{\mathrm{del}}$-step (a $\rightarrow_{\mathsf{S}}$-step), or the source of both a $\rightarrow_{@_0}$-step and a $\rightarrow_{@_1}$-step with the restriction that (in all three cases) it is not the source of any other step.

Heeding the fact that sources of $\rightarrow_\lambda$-steps are never sources of $\rightarrow_{@_i}$-steps, and vice versa, this condition can be relaxed to:

- no source of a $\rightarrow_\lambda$-step or a $\rightarrow_{@_i}$-step for $i \in \{0,1\}$ is also the source of a $\rightarrow_{\mathrm{del}}$-step ($\rightarrow_{\mathsf{S}}$-step), and every source of a $\rightarrow_{@_i}$-step for $i \in \{0,1\}$ is the source of both a $\rightarrow_{@_0}$-step and a $\rightarrow_{@_1}$-step, but not of any other step.

Definition 1.4.34 (eager scope/scope$^+$-delimiting strategy)**.** The *eager scope-delimiting strategy* $\mathbb{S}_{\mathrm{eag}}$ *for Reg* is defined as the restriction of rewrite steps in *Reg* to eager application of the rule ϱ^{del}: applications rules other than ϱ^{del} are

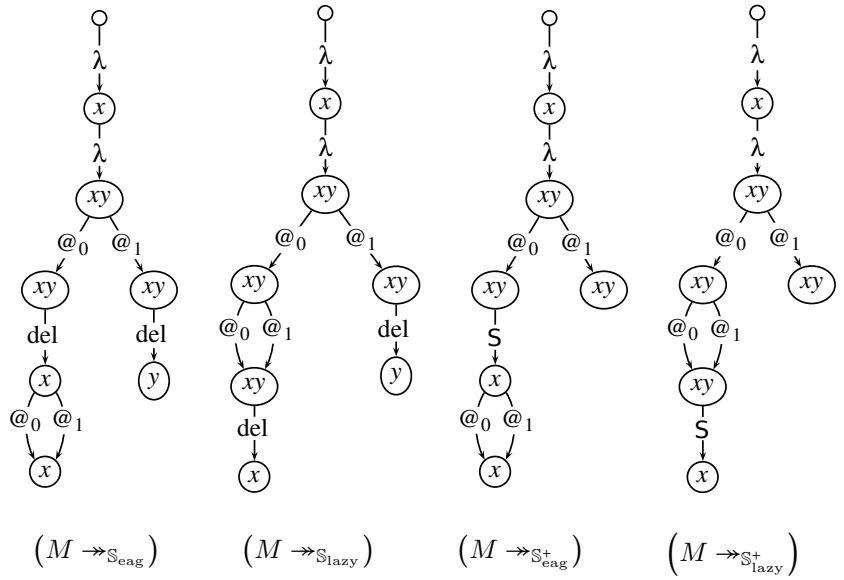

Figure 1.9. Reduction graphs of $M = \lambda x.\, \lambda y.\, x\ x\ y$ with respect to different *Reg* and *Reg*$^+$ strategies; compare: fig. 1.8. Again note, that the labels do only show the prefixes associated with each term.

only allowed if ϱ^{del} is not applicable. Analogously, the *eager scope*$^+$*-delimiting strategy* $\mathbb{S}^{+}_{\mathrm{eag}}$ *for Reg*$^+$ is defined as the restriction of rewrite steps in *Reg*$^+$ to eager application of the rule ϱ^{S}.

Definition 1.4.35 (lazy scope/scope$^+$-delimiting strategy)**.** The *lazy scope-delimiting strategy* $\mathbb{S}_{\mathrm{lazy}}$ *for Reg* is defined as the restriction of rewrite steps in *Reg* to lazy application of the rule ϱ^{del}: applications of ϱ^{del} are only allowed when other rules are not applicable. Analogously, the *lazy scope*$^+$*-delimiting strategy* $\mathbb{S}^{+}_{\mathrm{lazy}}$ *for Reg*$^+$ is defined as the restriction of rewrite steps in *Reg*$^+$ to lazy application of the rule ϱ^{S}.

§ 1.4.36 (history-aware versus history-free scope-delimiting strategies)**.** The history-free strategy obtained by projection from a history-aware scope-delimiting strategy is not in general a scope-delimiting strategy. This is due to the non-determinism which may be introduced by the projection. Consider

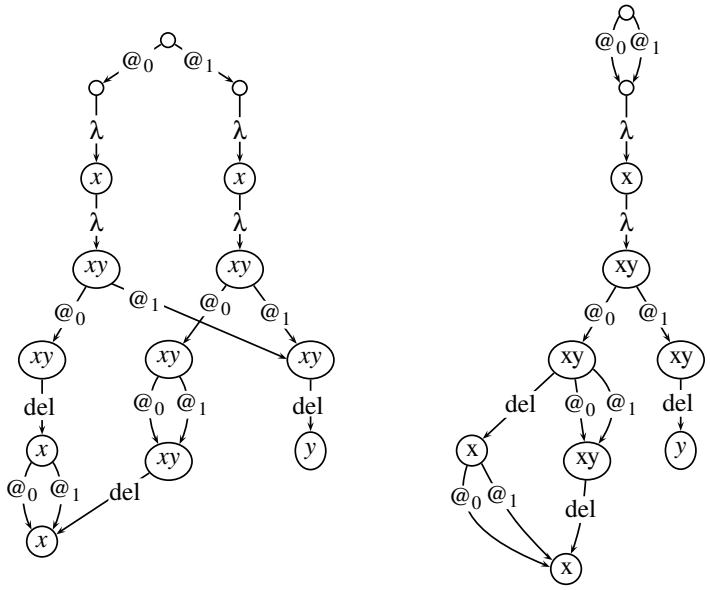

Figure 1.10. On the left: reduction graph of $M\ M$ with $M = \lambda x.\, \lambda y.\, x\ x\ y$ with respect to the history-aware strategy for *Reg* constructed by using $\mathbb{S}_{\text{eag}}$ on the left component and $\mathbb{S}_{\text{lazy}}$ on the right component of $M\ M$. On the right: reduction graph with respect to the history-free strategy obtained by projection. Note: vertex labels only show prefixes.

for example the term $M\ M$ with $M = \lambda x.\, \lambda y.\, x\ x\ y$ and the history-aware strategy constructed by using $\mathbb{S}_{\text{eag}}$ on the left component and $\mathbb{S}_{\text{lazy}}$ on the right component of $M\ M$. See fig. 1.10 for the reduction graph. The reduction graph w.r.t. the history-free strategy obtained by projection, however, bears the kind of non-determinism that is not permitted for a scope-delimiting strategy, in the form of a vertex that is the source of both a $\rightarrow_{\text{del}}$- and a $\rightarrow_{@_i}$-step.

The following proposition formulates a property of the eager scope-delimiting (scope$^+$-delimiting) strategy for *Reg* (in *Reg*$^+$) that assigns it a special status: the target of every rewrite sequence with respect to $\mathbb{S}_{\text{eag}}$ (with respect to $\mathbb{S}^+_{\text{eag}}$) can be reached, modulo some final $\rightarrow_{\text{del}}$-steps ($\rightarrow_{\mathsf{S}}$-steps), also by a rewrite sequence with respect to an arbitrary scope-delimiting (scope$^+$-delimiting) strategy. Furthermore, rewrite sequences with respect to $\mathbb{S}_{\text{eag}}$ (with respect to

$\mathbb{S}^{+}_{\mathrm{eag}}$) are able to mimic rewrite sequences with respect to an arbitrary scope-delimiting (scope^{+}-delimiting) strategy, up to trailing $\to_{\mathrm{del}}$-steps ($\to_{\mathsf{S}}$-steps) applied to the latter.

Definition 1.4.37. Let $\to_1$, $\to_2$, $\to_3$ be rewrite relations. The rewrite relation $\to_1$ is called *cofinal for* $\to_2$ if $\twoheadrightarrow_2 \subseteq \twoheadrightarrow_1 \cdot \twoheadleftarrow_2$. We say that $\to_1$ is *cofinal for* $\to_2$ *with trailing* $\to_3$*-steps* if $\twoheadrightarrow_2 \subseteq \twoheadrightarrow_1 \cdot \twoheadleftarrow_3$. Furthermore we say that $\to_1$ *factors into* $\to_2$ and $\to_3$ if $\to_1 \subseteq \to_2 \cdot \to_3$.

Proposition 1.4.38. For all scope-delimiting strategies $\mathbb{S}$ on *Reg* and all scope^{+}-delimiting strategies $\mathbb{S}^{+}$ on *Reg*$^{+}$ the following holds:

(i) $\twoheadrightarrow_{\mathbb{S}_{\mathrm{eag}}}$ factors into $\twoheadrightarrow_{\mathbb{S}}$ and $\twoheadrightarrow_{\mathrm{del}}$.
$\twoheadrightarrow_{\mathbb{S}^{+}_{\mathrm{eag}}}$ factors into $\twoheadrightarrow_{\mathbb{S}^{+}}$ and $\twoheadrightarrow_{\mathsf{S}}$.

(ii) $\to_{\mathbb{S}_{\mathrm{eag}}}$ is cofinal for $\to_{\mathbb{S}}$ with trailing $\to_{\mathrm{del}}$-steps.
$\to_{\mathbb{S}^{+}_{\mathrm{eag}}}$ is cofinal for $\to_{\mathbb{S}^{+}}$ with trailing $\to_{\mathsf{S}}$-steps.

Definition 1.4.39 (generated subterms)**.** Let $\mathbb{S}$ be a scope-delimiting strategy for *Reg*/*Reg*$^{+}$. For every $M \in \mathrm{Ter}^{\infty}(\boldsymbol{\lambda})$, the terms in the set $\mathrm{ST}_{\mathbb{S}}(M)$ are called the *generated subterms of* M *with respect to* $\mathbb{S}$, where $\mathrm{ST}_{\mathbb{S}}(M)$ is the set of objects of the generated sub-ARS $(() \; M \twoheadrightarrow_{\mathbb{S}})$, or in other words, the set of $\twoheadrightarrow_{\mathbb{S}}$-reducts of $() \; M$:

$$\begin{aligned} \mathrm{ST}_{\mathbb{S}} \; : \; \mathrm{Ter}^{\infty}(\boldsymbol{\lambda}) \; &\to \; \wp(\mathrm{Ter}^{\infty}((\boldsymbol{\lambda}))) \\ M \; &\mapsto \; O \quad \text{where } \langle O, \Phi, \mathrm{src}, \mathrm{tgt}\rangle = (() \; M \twoheadrightarrow_{\mathbb{S}}) \end{aligned}$$

Definition 1.4.40 ($\mathbb{S}$-regularity)**.** Let $\mathcal{A}$ be an abstract rewriting system, and $\mathbb{S}$ a strategy for $\mathcal{A}$. We say that an object o in $\mathcal{A}$ is $\mathbb{S}$*-regular in* $\mathcal{A}$ if the set of generated subterms $\mathrm{ST}_{\mathbb{S}}(o)$ of o is finite.

Definition 1.4.41 (regular and strongly regular λ-terms)**.** A λ-term M is called *regular* (*strongly regular*) if there exists a scope-delimiting strategy $\mathbb{S}$ for *Reg* (a scope^{+}-delimiting strategy $\mathbb{S}^{+}$ for *Reg*$^{+}$) such that M is $\mathbb{S}$-regular (is $\mathbb{S}^{+}$-regular).

Corollary 1.4.42 (regularity and strong regularity)**.** A λ-term M is regular (strongly regular) if and only if the set $\mathrm{ST}_{\mathbb{S}}(M)$ of generated subterms of M with respect to some scope-delimiting strategy (scope^{+}-delimiting strategy) $\mathbb{S}$ is finite.

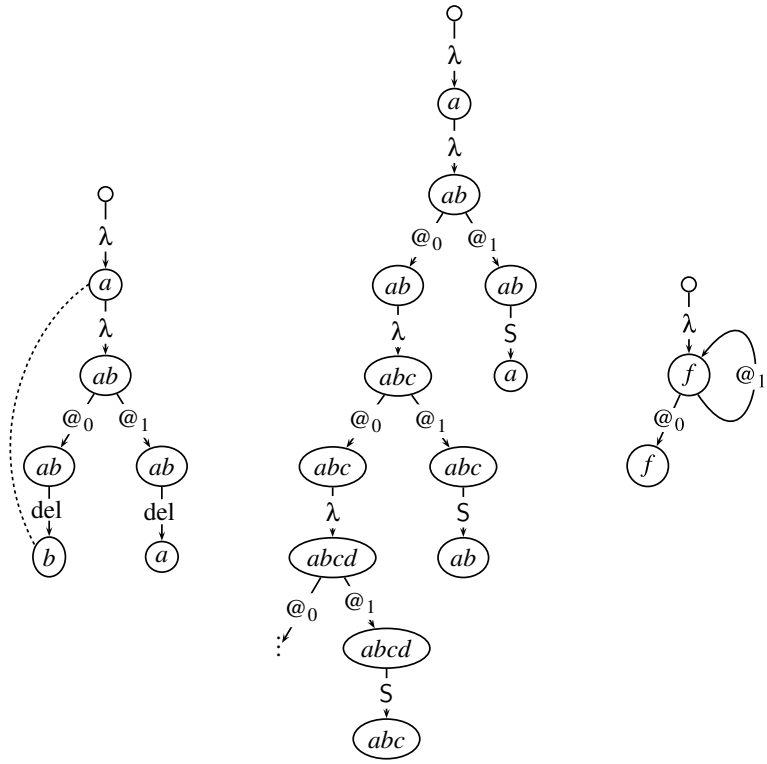

Figure 1.11. Reduction graphs (with prefixes as vertex labels) of:

left: example 1.2.2 w.r.t. $\mathbb{S}_{\text{eag}}$. The dotted line denotes node equality with the connected nodes representing identical (α-equivalent) terms. They are drawn as two separate nodes instead of one to avoid confusion as to which variable (it is a) is removed from the prefix by the incoming del-edge.

middle: example 1.2.2 w.r.t. an arbitrary scope$^+$-delimiting strategy for Reg^+ (and also with respect to $\mathbb{S}_{\text{lazy}}$). The vertical dots denote an infinite growth of the graph.

right: example 1.2.5 w.r.t. an arbitrary scope/scope$^+$-delimiting strategy for Reg^+/Reg.

Proposition 1.4.43 (finiteness ⇒ strong regularity ⇒ regularity)**.**

(i) Every strongly regular λ-term is also regular.

(ii) Finite λ-terms are strongly regular.

Proof. For statement (i), let M be a λ-term, and let $\mathbb{S}^+$ be a scope$^+$-delimiting strategy for Reg^+ by which $\mathrm{ST}_{\mathbb{S}^+}(M)$ is finite. Due to proposition 1.4.19 $\mathbb{S}^+$ can be modified to yield a scope-delimiting strategy $\mathbb{S}$ for Reg by replacing $\to_{\mathsf{S}}$-steps with $\to_{\mathrm{del}}$-steps. Then there is a stepwise correspondence between $\to_{\mathbb{S}^+}$-rewrite-sequences and $\to_{\mathbb{S}}$-rewrite-sequences that pass through the same terms. Consequently, the sets of $\twoheadrightarrow_{\mathbb{S}^+}$-reducts and $\twoheadrightarrow_{\mathbb{S}}$-reducts of $()\,M$ coincide: $\mathrm{ST}_{\mathbb{S}}(M) = \mathrm{ST}_{\mathbb{S}^+}(M)$. It follows that $\mathrm{ST}_{\mathbb{S}}(M)$ is finite.

For statement (ii), note that by proposition 1.4.24 (viii), and Kőnig's Lemma every finite term in $\mathrm{Ter}^\infty(\mathbf{(\lambda)})$ has only finitely many reducts with respect to $\twoheadrightarrow_{\mathrm{reg}}$ and $\twoheadrightarrow_{\mathrm{reg}^+}$. It follows that for every finite λ-term M and every scope-delimiting strategy $\mathbb{S}$ on Reg, or on Reg^+, the number of $\twoheadrightarrow_{\mathbb{S}}$-reducts of $()\,M$ is finite. □

Proposition 1.4.44 (eager scope-closure and regularity)**.** For all λ-terms M the following statements hold:

(i) M is regular if and only if M is $\mathbb{S}_{\mathrm{eag}}$-regular.

(ii) M is strongly regular if and only if M is $\mathbb{S}^+_{\mathrm{eag}}$-regular.

Proof. We only prove (i), because (ii) can be established analogously. The implication "⇐" follows from the definition of regularity. For "⇒", let M be a regular λ-term. Then there exists a scope-delimiting strategy $\mathbb{S}$ so that $\mathrm{ST}_{\mathbb{S}}(M)$ is finite. Since by proposition 1.4.38 (i) every $\twoheadrightarrow_{\mathbb{S}_{\mathrm{eag}}}$-rewrite-sequence factors into an $(\twoheadrightarrow_{\mathbb{S}} \cdot \twoheadrightarrow_{\mathrm{del}})$-rewrite-sequence, it follows that every term in $\mathrm{ST}_{\mathbb{S}_{\mathrm{eag}}}(M)$ is the $\twoheadrightarrow_{\mathrm{del}}$-reduct of a term in $\mathrm{ST}_{\mathbb{S}}(M)$. As every term in $\mathrm{Ter}^\infty(\mathbf{(\lambda)})$ has only finitely many $\twoheadrightarrow_{\mathrm{del}}$-reducts, it follows that also $\mathrm{ST}_{\mathbb{S}_{\mathrm{eag}}}(M)$ is finite. □

Example 1.4.45 (regular and strongly regular terms)**.** The following examples demonstrate the connection between (strong) regularity and, as illustrated in fig. 1.11, the finiteness of the ARSs generated by different Reg/Reg^+ strategies.

- example 1.2.2 is regular but not strongly regular.
- example 1.2.5 is strongly regular.

To further illustrate the statements made in example 1.4.45 let us consider various *Reg* and Reg^+ rewrite sequences corresponding to infinite paths through the terms.

Example 1.4.46. For the term M from example 1.2.2, we first introduce a finite CRS based notation, as a 'higher-order recursive program scheme'. We can represent M by $\lambda a.\, \mathsf{rec}_M(a)$ together with the CRS rule $\mathsf{rec}_M(X) \to_{\mathsf{rec}_M} \lambda x.\, \mathsf{rec}_M(x)\ X$. It holds that $\lambda a.\, \mathsf{rec}_M(a) \twoheadrightarrow_{\mathsf{rec}_M} M$. Using this notation we can finitely describe the infinite path down the spine of example 1.2.2 by the cyclic $\mathbb{S}_{\mathrm{eag}}$-rewrite-sequence:

$$
\begin{array}{ll}
 & () \ \lambda a.\, \mathsf{rec}_M(a) \\
\to_{\lambda} & (\lambda a)\ \mathsf{rec}_M(a) \\
\to_{\mathsf{rec}_M} & (\lambda a)\ \lambda b.\, \mathsf{rec}_M(b)\ a \\
\to_{\lambda} & (\lambda ab)\ \mathsf{rec}_M(b)\ a \\
\to_{@_0} & (\lambda ab)\ \mathsf{rec}_M(b) \\
\to_{\mathrm{del}} & (\lambda b)\ \mathsf{rec}_M(b) \\
= & (\lambda a)\ \mathsf{rec}_M(a)
\end{array}
$$

In Reg^+ the rewriting sequence for the same path is invariant over all scope$^+$-delimiting strategies and necessarily infinite:

$$
\begin{array}{llll}
() \ \lambda a.\, \mathsf{rec}_M(a) & \to_{\lambda} (\lambda a)\ \mathsf{rec}_M(a) & \to_{\mathsf{rec}_M} & \\
(\lambda a)\ \lambda b.\, \mathsf{rec}_M(b)\ a & \to_{\lambda} (\lambda ab)\ \mathsf{rec}_M(b)\ a & \to_{@_0} (\lambda ab)\ \mathsf{rec}_M(b) & \to_{\mathsf{rec}_M} \\
(\lambda ab)\ \lambda c.\, \mathsf{rec}_M(c)\ b & \to_{\lambda} (\lambda abc)\ \mathsf{rec}_M(c)\ b & \to_{@_0} (\lambda abc)\ \mathsf{rec}_M(c) & \to_{\mathsf{rec}_M} \\
(\lambda abc)\ \lambda d.\, \mathsf{rec}_M(d)\ c & \to_{\lambda} (\lambda abcd)\ \mathsf{rec}_M(d)\ c & \to_{@_0} \ldots &
\end{array}
$$

Example 1.4.47. As an illustration of a regular term we study M defined as the unfolding of $\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f$ from example 1.2.4. It is strongly regular since the infinite path through the term can be witnessed by this cyclic Reg^+

rewriting sequence:

$$\begin{array}{ll} & () \; M \\ = & () \; \lambda xy.\, M \; y \; x \\ \to_{\lambda} & (\lambda x) \; \lambda y.\, M \; y \; x \\ \to_{\lambda} & (\lambda xy) \; M \; y \; x \\ \to_{@_0} & (\lambda xy) \; M \; y \\ \to_{@_0} & (\lambda xy) \; M \\ \to_{\mathsf{S}} & (\lambda x) \; M \\ \to_{\mathsf{S}} & () \; M \\ & \dots \end{array}$$

See also fig. 1.12 for a graphical illustration of the reduction graph.

§ 1.4.48 (eager scope-closure is necessary)**.** The restriction of proposition 1.4.44 to the eager scope-delimiting strategy cannot be relaxed to arbitrary scope-delimiting strategies. The term in example 1.4.47 for instance is $\mathbb{S}^{+}_{\text{eag}}$-regular but not $\mathbb{S}^{+}_{\text{lazy}}$-regular (see fig. 1.12).

Definition 1.4.49 (grounded cycles in *Reg*, *Reg*$^{+}$)**.** Let $\tau : (\lambda\vec{x}_0) \; M_0 \to (\lambda\vec{x}_1) \; M_1 \to \dots$ be a finite or infinite rewrite sequence with respect to $\to_{\text{reg}}$ or $\to_{\text{reg}^+}$. By a *grounded cycle* in τ we mean a cycle $(\lambda\vec{x}_i) \; M_i \to (\lambda\vec{x}_{i+1}) \; M_{i+1} \to \dots \to (\lambda\vec{x}_{i+k}) \; M_{i+k} = (\lambda\vec{x}_i) \; M_i$, in τ, where $i \in \mathbb{N}$ and $k \geq 1$, with the additional property that $|\vec{x}_{i+j}| \geq |\vec{x}_i|$ for all $j \in \{0, \dots, k\}$ (i.e. the lengths of the abstraction prefixes in the terms of the cycle is greater or equal to the length of the abstraction prefix at the first and final term of the cycle).

Proposition 1.4.50 (infinite $\mathbb{S}$-rewrite-sequences contain grounded cycles)**.** Let M be a λ-term that is $\mathbb{S}$-regular for some scope/scope^{+}-delimiting strategy $\mathbb{S}$. Then every infinite rewrite sequence with respect to $\mathbb{S}$ contains a grounded cycle.

Proof. Since the argument is analogous in both cases, we only treat strongly regular terms. Let M be a λ-term that is $\mathbb{S}^{+}$-regular for some scope^{+}-delimiting strategy $\mathbb{S}^{+}$, and let $\tau : M = (\lambda\vec{x}_0) \; M_0 \to_{\mathbb{S}^{+}} (\lambda\vec{x}_1) \; M_1 \to_{\mathbb{S}^{+}} \dots$ be an infinite rewrite sequence. As M is $\mathbb{S}^{+}$-regular, the sequence $\{(\lambda\vec{x}_i) \; M_i\}_{i\in\mathbb{N}}$ of terms on τ contains only finitely many different terms. Let $l := \liminf\{|\vec{x}_i|\}_{i\in\mathbb{N}}$, that is,

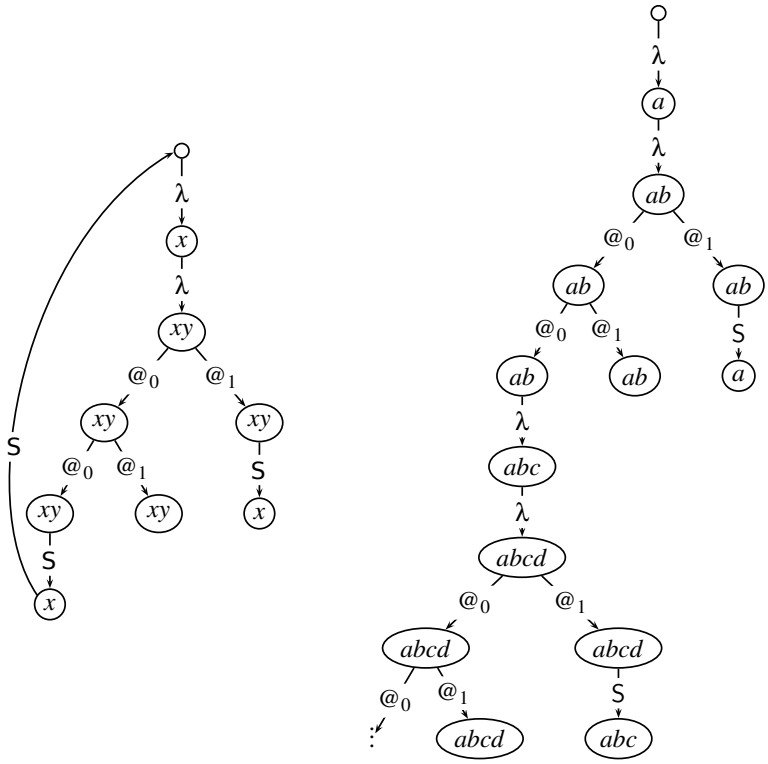

Figure 1.12. ***Reg***$^+$-reduction-graphs of $[\![\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f]\!]_{\boldsymbol{\lambda}}$ with respect to the scope$^+$-delimiting strategies $\mathbb{S}^+_{\mathrm{eag}}$ (left) and $\mathbb{S}^+_{\mathrm{lazy}}$ (right).

the minimum of abstraction prefix lengths that appears infinitely often on τ. Let $\{(\lambda\vec{x}_{i_j})\ M_{i_j}\}_{j\in\mathbb{N}}$ be the subsequence of $\{(\lambda\vec{x}_i)\ M_i\}_{i\in\mathbb{N}}$ consisting of terms with prefix length l, and such that, for all $k \geq i_0$, $|\vec{x}_k| \geq l$. Since also this subsequence contains only finitely many terms, there exist $j_1, j_2 \in \mathbb{N}$, $j_1 < j_2$ such that $(\lambda\vec{x}_{i_{j_1}})\ M_{i_{j_1}} = (\lambda\vec{x}_{i_{j_2}})\ M_{i_{j_2}}$. By the choice of the subsequence it follows that $(\lambda\vec{x}_{i_{j_1}})\ M_{i_{j_1}} \to_{\mathbb{S}^+} \ldots \to_{\mathbb{S}^+} (\lambda\vec{x}_{i_{j_2}})\ M_{i_{j_2}}$ is a grounded cycle in τ. □

§ 1.4.51. We round off this section by providing a motivation for the system Reg^+ in terms of an operation 'parse' that when applied to a λ-term (i) decomposes it into its generated subterms, and (ii) recombines the generated subterms encountered in the decomposition analysis with, in the limit, the original term as the result. With this purpose in mind, we define a CRS $\boldsymbol{Parse^+}$.

Definition 1.4.52 ($\boldsymbol{Parse^+}$)**.** Let $\Sigma^{(\lambda)}_{\mathrm{parse}^+} = \Sigma^{(\lambda)} \cup \{\mathsf{parse}^+_n \mid n \in \mathbb{N}\}$ be the extension of the CRS signature for $\boldsymbol{(\lambda)}$, where for $n \in \mathbb{N}$, the symbols parse^+_n have arity n. By $\boldsymbol{Parse^+}$ we denote the CRS over $\Sigma^{(\lambda)}_{\mathrm{parse}^+}$ with the following rules:

$$\varrho^{@}_{\mathrm{parse}^+}\colon\ \mathsf{parse}^+_n\big(\vec{X}_n, \mathsf{pre}_n([\vec{x}_n]\,\mathsf{app}(Z_0(\vec{x}_n),\, Z_1(\vec{x}_n)))\big) \to \mathsf{app}\big(\mathsf{parse}^+_n\big(\vec{X}_n, \mathsf{pre}_n([\vec{x}_n]\,Z_0(\vec{x}_n))\big),\ \mathsf{parse}^+_n\big(\vec{X}_n, \mathsf{pre}_n([\vec{x}_n]\,Z_1(\vec{x}_n))\big)\big)$$

$$\varrho^{\lambda}_{\mathrm{parse}^+}\colon\ \mathsf{parse}^+_n\big(\vec{X}_n, \mathsf{pre}_n([\vec{x}_n]\,\mathsf{abs}([y]\,Z(\vec{x}_n, y)))\big) \to \mathsf{abs}\big([y]\,\mathsf{parse}^+_{n+1}\big(\vec{X}_n, y, \mathsf{pre}_{n+1}([\vec{x}_n]\,[y]\,Z(\vec{x}_n, y))\big)\big)$$

$$\varrho^{\mathsf{S}}_{\mathrm{parse}^+}\colon\ \mathsf{parse}^+_{n+1}()X_1, \ldots, X_{n+1}(\mathsf{pre}_{n+1}([\vec{x}_{n+1}]\,Z(\vec{x}_n)) \to \mathsf{parse}^+_n(X_1, \ldots, X_n, \mathsf{pre}_n([\vec{x}_n]\,Z(\vec{x}_n)))$$

$$\varrho^{0}_{\mathrm{parse}^+}\colon\ \mathsf{parse}^+_n(X_1, \ldots, X_n, \mathsf{pre}_n([\vec{x}_n]\,x_n)) \to X_n$$

We denote by $\to_{\mathrm{parse}^+}$ the rewrite relation induced by this CRS.

§ 1.4.53 ($\boldsymbol{Parse^+}$ contains $\boldsymbol{Reg^+}$)**.** Observe that the rules $\varrho^{@_i}$ for $i \in \{0, 1\}$, ϱ^{λ}, and ϱ^{S} of the CRS $\boldsymbol{Reg^+}$ are contained within the rules $\varrho^{@}_{\mathrm{parse}^+}$, $\varrho^{\lambda}_{\mathrm{parse}^+}$, $\varrho^{\mathsf{S}}_{\mathrm{parse}^+}$, respectively, of the CRS $\boldsymbol{Parse^+}$, in the sense that applications of the latter rules include applications of the former. This has as a consequence that repeated $\to_{\mathrm{parse}^+}$-steps on a term $()\ M$ lead to terms that contain generated subterms of M as closed subexpressions. Furthermore $\to_{\mathrm{parse}^+}$-rewrite-sequences on $()\ M$ are possible that move redexes simultaneously deeper and deeper, analysing

ever larger parts of M, and at the same time recreating a larger and larger λ-term parts (stable prefix contexts) of M, the original term.

Proposition 1.4.54. For every term $M \in \mathrm{Ter}^{\infty}(\boldsymbol{\lambda})$ it holds:

(i) ***Parse***$^{+}$ analyses M into its generated subterms: If $\mathsf{parse}_0^+(()\ M) \twoheadrightarrow_{\mathrm{parse}^+} M'$, then all subexpressions starting with pre_n (for some $n \in \mathbb{N}$) in M' are generated subterms of M. Moreover, for every generated subterm $(\lambda\vec{y})\ N$ of M, there exists a $\twoheadrightarrow_{\mathrm{parse}^+}$-reduction M'' of M such that $(\lambda\vec{y})\ N$ is contained in M''.

(ii) ***Parse***$^{+}$ reconstructs M: $\mathsf{parse}_0^+(()\ M) \twoheadrightarrow\!\!\!\!\twoheadrightarrow_{\mathrm{parse}^+} M$.

Example 1.4.55. Let M be the infinite unfolding of $\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f$ for which we use as a finite representation the equation $M = \lambda xy.\, M\ y\ x$. In ***Parse***$^{+}$, M is decomposed, and composed again, by the infinite rewrite sequence (see also example 1.2.4):

$$
\begin{array}{ll}
 & \mathsf{parse}_0^+(()\ M) \\
= & \mathsf{parse}_0^+(()\ \lambda x.\, \lambda y.\, M\ y\ x) \\
\to_{\mathrm{parse}^+.\lambda} & \lambda x'.\, \mathsf{parse}_1^+(x', (\lambda x')\ \lambda y.\, M\ y\ x') \\
\to_{\mathrm{parse}^+.\lambda} & \lambda x'.\, \lambda y'.\, \mathsf{parse}_2^+(x', y', (\lambda x' y')\ M\ y'\ x') \\
\to_{\mathrm{parse}^+.@} & \lambda x'.\, \lambda y'.\, \mathsf{parse}_2^+(x', y', (\lambda x' y')\ M\ y')\ \mathsf{parse}_2^+(x', y', (\lambda x' y')\ x') \\
\to_{\mathrm{parse}^+.\mathsf{S}} & \lambda x'.\, \lambda y'.\, \mathsf{parse}_2^+(x', y', (\lambda x' y')\ M\ y')\ \mathsf{parse}_1^+(x', (\lambda x')\ x') \\
\to_{\mathrm{parse}^+.0} & \lambda x'.\, \lambda y'.\, \mathsf{parse}_2^+(x', y', (\lambda x' y')\ M\ y')\ x' \\
\to_{\mathrm{parse}^+.@} & \lambda x'.\, \lambda y'.\, \mathsf{parse}_2^+(x', y', (\lambda x' y')\ M)\ \mathsf{parse}_2^+(x', y', (\lambda x' y')\ y')\ x' \\
\to_{\mathrm{parse}^+.0} & \lambda x'.\, \lambda y'.\, \mathsf{parse}_2^+(x', y', (\lambda x' y')\ M)\ y'\ x' \\
\to_{\mathrm{parse}^+.\mathsf{S}} & \lambda x'.\, \lambda y'.\, \mathsf{parse}_1^+(x', (\lambda x')\ M)\ y'\ x' \\
\to_{\mathrm{parse}^+.\mathsf{S}} & \lambda x'.\, \lambda y'.\, \mathsf{parse}_0^+(()\ M)\ y'\ x' \\
= & \lambda x'.\, \lambda y'.\, \mathsf{parse}_0^+(()\ \lambda x.\, \lambda y.\, M\ y\ x)\ y'\ x' \\
\to_{\mathrm{parse}^+.\lambda} & \cdots
\end{array}
$$

Note that the generated subterms of M appear as the last arguments of the parse_i^+-symbols in this rewrite sequence.

1.5 Observing $\lambda_{\mathsf{letrec}}$-terms by their generated subterms

§ 1.5.1 (overview)**.** In this section we adapt the concepts developed so far for the infinitary λ-calculus to $\lambda_{\mathsf{letrec}}$. By combining the rules of $\boldsymbol{R}_{\triangledown}$ with those of $\boldsymbol{Reg}$ and $\boldsymbol{Reg}^{+}$, respectively, we obtain the CRSs $\boldsymbol{Reg}_{\mathsf{letrec}}$ and $\boldsymbol{Reg}^{+}_{\mathsf{letrec}}$ for the deconstruction of $\lambda_{\mathsf{letrec}}$-terms furnished with an abstraction prefix. We define scope-delimiting and scope^{+}-delimiting strategies for $\boldsymbol{Reg}_{\mathsf{letrec}}$ and $\boldsymbol{Reg}^{+}_{\mathsf{letrec}}$ as before by excluding all non-determinism except for sources of $\to_{@_0}$-steps and $\to_{@_1}$-steps.

Definition 1.5.2 (the CRSs $\boldsymbol{Reg}_{\mathsf{letrec}}$ and $\boldsymbol{Reg}^{+}_{\mathsf{letrec}}$ for decomposing $\lambda_{\mathsf{letrec}}$-terms)**.** We extend $\Sigma^{\lambda}_{\mathsf{letrec}}$ (see definition 0.5.4) by function symbols pre_n with arity one to obtain the signature $\Sigma^{(\lambda)}_{\mathsf{letrec}} = \Sigma^{\lambda}_{\mathsf{letrec}} \cup \{\mathsf{pre}_n \mid n \in \mathbb{N}\}$. We denote the induced set of (finite) prefixed $\lambda_{\mathsf{letrec}}$-terms by $\mathrm{Ter}(\boldsymbol{(\lambda_{\mathsf{letrec}})})$ and we adopt the same informal notation for $\boldsymbol{(\lambda_{\mathsf{letrec}})}$ as for $\boldsymbol{(\lambda)}$ (see definition 1.4.13). On the signature $\Sigma^{\lambda}_{\mathsf{letrec}}$ we define the CRS $\boldsymbol{Reg}_{\mathsf{letrec}}$ (the CRS $\boldsymbol{Reg}^{+}_{\mathsf{letrec}}$) with the rules as the union of the rules of $\boldsymbol{R}_{\triangledown}$ and $\boldsymbol{Reg}$ (the union of the rules of $\boldsymbol{R}_{\triangledown}$ and $\boldsymbol{Reg}^{+}$).

§ 1.5.3 (list of rules for $\boldsymbol{Reg}_{\mathsf{letrec}}/\boldsymbol{Reg}^{+}_{\mathsf{letrec}}$)**.** The rules of $\boldsymbol{Reg}_{\mathsf{letrec}}/\boldsymbol{Reg}^{+}_{\mathsf{letrec}}$ consist of the set of rules arising from the rule schemes $\varrho^{\lambda}_{\triangledown}$, $\varrho^{@}_{\triangledown}$, $\varrho^{\mathrm{letrec}}_{\triangledown}$, $\varrho^{\mathrm{rec}}_{\triangledown}$, $\varrho^{\mathrm{nil}}_{\triangledown}$, $\varrho^{\mathrm{red}}_{\triangledown}$ from the unfolding CRS $\boldsymbol{R}_{\triangledown}$, joined with the rules arising from the rule schemes $\varrho^{@_0}_{\Delta}$, $\varrho^{@_1}_{\Delta}$, $\varrho^{\lambda}_{\Delta}$, $\varrho^{\mathrm{del}}_{\Delta}/\varrho^{\mathsf{S}}_{\Delta}$ from the decomposition CRS $\boldsymbol{Reg}/\boldsymbol{Reg}^{+}$.

Notation 1.5.4 (rewrite relations for $\boldsymbol{Reg}_{\mathsf{letrec}}/\boldsymbol{Reg}^{+}_{\mathsf{letrec}}$)**.** To avoid ambiguity when referring to the rewriting relations induced by the rules of $\boldsymbol{Reg}_{\mathsf{letrec}}/\boldsymbol{Reg}^{+}_{\mathsf{letrec}}$, we will prefix the rule names of $\boldsymbol{R}_{\triangledown}$ with $\triangledown$. Thus, we will denote the induced rewrite relations for $\boldsymbol{Reg}_{\mathsf{letrec}}/\boldsymbol{Reg}^{+}_{\mathsf{letrec}}$ as: $\to_{\triangledown.\lambda}$, $\to_{\triangledown.@}$, $\to_{\triangledown.\mathrm{letrec}}$, $\to_{\triangledown.\mathrm{rec}}$, $\to_{\triangledown.\mathrm{nil}}$, $\to_{\triangledown.\mathrm{red}}$, $\to_{@_0}$, $\to_{@_1}$, $\to_{\lambda}$, $\to_{\mathrm{del}}/\to_{\mathsf{S}}$.

Definition 1.5.5 (the ARSs Reg_{letrec}, $Reg^{+}_{\mathsf{letrec}}$)**.** By Reg_{letrec}, and $Reg^{+}_{\mathsf{letrec}}$ we denote the ARSs that result by restricting the ARSs induced by $\boldsymbol{Reg}_{\mathsf{letrec}}$, and $\boldsymbol{Reg}^{+}_{\mathsf{letrec}}$, respectively, to the subset $\mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$ of terms.

The lemmas below state a number of simple rewrite properties for Reg_{letrec} and $Reg^{+}_{\mathsf{letrec}}$ concerning the interplay between unfolding and decomposition steps.

Lemma 1.5.6. On $\mathrm{Ter}(\boldsymbol{(\lambda_{\mathsf{letrec}})})$, the rewrite relations in $\boldsymbol{Reg}_{\mathsf{letrec}}$ and $\boldsymbol{Reg}^{+}_{\mathsf{letrec}}$ have the following properties:

(i) $\to_\triangledown$ one-step commutes with $\to_\lambda$, $\to_{@_0}$, $\to_{@_1}$, $\to_{\mathsf{S}}$, and $\to_{\mathrm{del}}$:

$$\begin{array}{lcll} \leftarrow_\triangledown \cdot \to_\lambda & \subseteq & \to_\lambda \cdot \leftarrow_\triangledown & \\ \leftarrow_\triangledown \cdot \to_{@_i} & \subseteq & \to_{@_i} \cdot \leftarrow_\triangledown & (i \in \{0,1\}) \\ \leftarrow_\triangledown \cdot \to_{\mathsf{S}} & \subseteq & \to_{\mathsf{S}} \cdot \leftarrow_\triangledown & \\ \leftarrow_\triangledown \cdot \to_{\mathrm{del}} & \subseteq & \to_{\mathrm{del}} \cdot \leftarrow_\triangledown & \end{array}$$

(ii) Reg_{letrec} and Reg^+_{letrec} have the same normal forms as Reg and Reg^+, respectively: $(\lambda x)\ x$ is the only term in $\mathrm{Ter}(\boldsymbol{(\lambda_{\mathsf{letrec}})})$ in $\to_{\mathrm{reg}}$-normal-form. Every $\to_{\mathrm{reg}^+}$-normal-form in $\mathrm{Ter}(\boldsymbol{(\lambda_{\mathsf{letrec}})})$ is of the form $(\lambda x_1 \ldots x_n)\ x_n$.

Proof. The commutation properties in (i) are easy to verify by analysing the behaviour of the rewrite rules in Reg_{letrec} and in Reg^+_{letrec} on the terms of $\mathrm{Ter}(\boldsymbol{(\lambda_{\mathsf{letrec}})})$.

The statement in (ii) follows from proposition 1.4.24 (i), and (viii): Normal forms with respect to Reg_{letrec} and Reg^+_{letrec} can only be λ-terms without occurrences of letrec, since every occurrence of letrec in a $\lambda_{\mathsf{letrec}}$-term gives rise to a $\to_\triangledown$-redex. □

Lemma 1.5.7. The rewrite relation $\to_\triangledown^{!\omega}$ (to $\to_\triangledown$-normal-form in at most ω steps) one-step commutes with $\to_\lambda$, $\to_{@_0}$, $\to_{@_1}$, $\to_{\mathsf{S}}$, and $\to_{\mathrm{del}}$:

$$\begin{array}{llll} \leftarrow_\triangledown^{!\omega} \cdot \to_\lambda \subseteq \to_\lambda \cdot \leftarrow_\triangledown^{!\omega} & & \leftarrow_\triangledown^{!\omega} \cdot \to_{@_i} \subseteq \to_{@_i} \cdot \leftarrow_\triangledown^{!\omega} & (i \in \{0,1\}) \\ \leftarrow_\triangledown^{!\omega} \cdot \to_{\mathsf{S}} \subseteq \to_{\mathrm{del}} \cdot \leftarrow_\triangledown^{!\omega} & & \leftarrow_\triangledown^{!\omega} \cdot \to_{\mathrm{del}} \subseteq \to_{\mathrm{del}} \cdot \leftarrow_\triangledown^{!\omega} & \end{array}$$

This implies, for prefixed terms that have unfoldings that:

$$\begin{array}{c} [\![(\lambda\vec{x})\ \lambda y.\, L_0]\!]_{\boldsymbol{\lambda}} \to_\lambda [\![(\lambda\vec{x}y)\ L_0]\!]_{\boldsymbol{\lambda}} \\ [\![(\lambda\vec{x})\ L_0\ L_1]\!]_{\boldsymbol{\lambda}} \to_{@_i} [\![(\lambda\vec{x})\ L_i]\!]_{\boldsymbol{\lambda}} \quad (i \in \{0,1\}) \\ (\lambda\vec{x})\ M \to_{\mathsf{S}} (\lambda\vec{x}')\ M \ \Rightarrow\ [\![(\lambda\vec{x})\ M]\!]_{\boldsymbol{\lambda}} \to_{\mathsf{S}} [\![(\lambda\vec{x}')\ M]\!]_{\boldsymbol{\lambda}} \\ (\lambda\vec{x})\ M \to_{\mathrm{del}} (\lambda\vec{x}')\ M \ \Rightarrow\ [\![(\lambda\vec{x})\ M]\!]_{\boldsymbol{\lambda}} \to_{\mathrm{del}} [\![(\lambda\vec{x}')\ M]\!]_{\boldsymbol{\lambda}} \end{array}$$

Furthermore it holds: $\leftarrow_\triangledown^{!\omega} \cdot \to_\triangledown \subseteq \leftarrow_\triangledown^{!\omega}$.

Proof. The commutation properties with the rewrite relation $\to_\triangledown^{!\omega}$ can be shown by using refined versions of the commutation properties in lemma 1.5.6 (i), in which the minimal depth of unfolding steps is taken account of. When denoting

by $\to_{\triangledown}^{\geq n}$ the rewrite relation that is generated by $\to_{\triangledown}$-steps of depth $\geq n$, then the following properties hold:

$$\leftarrow_{\triangledown}^{\geq n} \cdot \to_{\lambda} \subseteq \to_{\lambda} \cdot \leftarrow_{\triangledown}^{\geq n} \qquad \leftarrow_{\triangledown}^{\geq n+1} \cdot \to_{@_i} \subseteq \to_{@_i} \cdot \leftarrow_{\triangledown}^{\geq n}$$
$$\leftarrow_{\triangledown}^{\geq n+1} \cdot \to_{\mathsf{S}} \subseteq \to_{\mathsf{S}} \cdot \leftarrow_{\triangledown}^{\geq n} \qquad \leftarrow_{\triangledown}^{\geq n+1} \cdot \to_{\mathrm{del}} \subseteq \to_{\mathrm{del}} \cdot \leftarrow_{\triangledown}^{\geq n}$$

Using these properties, strongly convergent $\to_{\triangledown}$-rewrite-sequences can be shown to project, via $\to_{\lambda}$-, $\to_{@_i}$-, $\to_{\mathrm{del}}$-, and $\to_{\mathsf{S}}$-steps, to strongly convergent $\to_{\triangledown}$-rewrite-sequences.

The property $\leftarrow_{\triangledown}^{!\omega} \cdot \to_{\triangledown} \subseteq \leftarrow_{\triangledown}^{!\omega}$ can be shown by using refined versions of the elementary diagrams from the confluence proof that take the minimal depths of steps into account. □

§ 1.5.8 (scope/scope$^+$-delimiting strategy for $Reg_{\mathsf{letrec}}/Reg^+_{\mathsf{letrec}}$)**.** As for Reg/Reg^+ we require of scope/scope$^+$-delimiting strategies to have deterministic $\to_{\mathrm{del}}/\to_{\mathsf{S}}$-steps. As we did for definition 1.4.32, we will also fix all non-determinism except for the choice between $@_0$ and $@_1$.

Definition 1.5.9 (scope/scope$^+$-delimiting strategy for $Reg_{\mathsf{letrec}}/Reg^+_{\mathsf{letrec}}$)**.** A strategy $\mathbb{S}$ for Reg_{letrec} (Reg^+_{letrec}) will be called a *scope-delimiting* (*scope$^+$-delimiting*) *strategy* if:

- $\mathbb{S}$ is deterministic for sources of $\to_{\lambda}$-steps, $\to_{\mathrm{del}}$-steps ($\to_{\mathsf{S}}$-steps), and all letrec-unfolding steps (i.e. all $\to_{\triangledown.\rho}$-steps for every rule $\varrho^{\rho}_{\triangledown}$ of $\boldsymbol{R}_{\boldsymbol{\triangledown}}$).
- $\mathbb{S}$ enforces eager application of $\varrho^{\mathrm{red}}_{\triangledown}$: every source of a step in $\mathbb{S}$ according to an application of a rule different from $\varrho^{\mathrm{red}}_{\triangledown}$ is not the source of a $\to_{\mathrm{red}}$-step in the underlying ARS.

We say that such a strategy $\mathbb{S}$ is a *lazy-unfolding* scope-delimiting strategy (a *lazy-unfolding* scope$^+$-delimiting strategy) if furthermore:

- $\mathbb{S}$ applies the rules of $\boldsymbol{R}_{\boldsymbol{\triangledown}}$ except for $\varrho^{\mathrm{red}}_{\triangledown}$ only at the root of the term, i.e. directly beneath the abstraction prefix.
- $\mathbb{S}$ uses the rules of $\boldsymbol{R}_{\boldsymbol{\triangledown}}$ other than $\varrho^{\mathrm{red}}_{\triangledown}$ in a lazy way: every source of a step in $\mathbb{S}$ with respect to a rule of $\boldsymbol{R}_{\boldsymbol{\triangledown}}$ other than $\varrho^{\mathrm{red}}_{\triangledown}$ is not also the source of a step in the underlying ARS, with respect to one of the rules of *Reg* (of Reg^+).

Notation 1.5.10. For every scope-delimiting strategy $\mathbb{S}$ on Reg_{letrec} (on $Reg^{+}_{\mathsf{letrec}}$), we denote by $\to_{\mathbb{S}.\rho}$ the rewrite relation that is induced by those steps according to $\mathbb{S}$ due to applications of the rule ρ of Reg_{letrec} (of $Reg^{+}_{\mathsf{letrec}}$).

§ 1.5.11 (deterministic unfolding)**.** Note that in definition 1.5.9 we do not only require of a strategy to eliminate the non-determinism with respect to $\varrho^{\mathrm{del}}_{\Delta}$-steps ($\varrho^{\mathsf{S}}_{\Delta}$-steps) but all non-determinism except for the $\varrho^{@_0}_{\Delta}/\varrho^{@_1}_{\Delta}$-non-determinism. This restriction will play a role later for the definition of λ-transition-graphs in section 1.9, and aso in this section for defining projections of scope/scope^{+}-delimiting strategies for $\lambda_{\mathsf{letrec}}$-terms to scope/scope^{+}-delimiting strategies for λ-terms.

§ 1.5.12 (eager application of $\varrho^{\mathrm{red}}_{\triangledown}$)**.** By requiring scope/scope^{+} delimiting strategies to apply $\varrho^{\mathrm{red}}_{\triangledown}$ eagerly we can exploit a useful property with respect to free variables of a term: if $(\lambda\vec{x})\ L \in \mathrm{Ter}(\boldsymbol{(\lambda_{\mathsf{letrec}})})$ is in $\to_{\triangledown.\mathrm{red}}$-normal-form, then the free variables occurring in M correspond to the free variables of $[\![L]\!]_{\boldsymbol{\lambda}}$.

§ 1.5.13 (unfolding of prefixed terms)**.** Note that we have just applied $[\![\cdot]\!]_{\boldsymbol{\lambda}}$ to a prefixed term while in definition 0.7.1 $[\![\cdot]\!]_{\boldsymbol{\lambda}}$ is only defined for 'normal', unprefixed terms from $\mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$. We will continue to do so in the rest of this section and assume that the domain of $[\![\cdot]\!]_{\boldsymbol{\lambda}}$ is extended in the obvious way to terms in $\mathrm{Ter}(\boldsymbol{(\lambda_{\mathsf{letrec}})})$.

Definition 1.5.14 ($\mathbb{S}$-productive terms)**.** Let L be a $\lambda_{\mathsf{letrec}}$-term, and $\mathbb{S}$ a scope-delimiting strategy for Reg_{letrec} or a scope^{+}-delimiting strategy for $Reg^{+}_{\mathsf{letrec}}$. We say that L is $\mathbb{S}$-*productive* if every infinite rewrite sequence on L with respect to $\mathbb{S}$ contains infinitely many steps according to $\to_{\mathbb{S}.@_0}$, $\to_{\mathbb{S}.@_1}$, or $\to_{\mathbb{S}.\lambda}$.

Definition 1.5.15 (generated subterms of $\lambda_{\mathsf{letrec}}$-terms)**.** We extend the domain of ST from definition 1.4.39 to $\lambda_{\mathsf{letrec}}$-terms. Let $\mathbb{S}$ be a scope-delimiting strategy for $Reg_{\mathsf{letrec}}/Reg^{+}_{\mathsf{letrec}}$. For every $L \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$, $\mathrm{ST}_{\mathbb{S}}(L)$ is defined as the set of objects of the sub-ARS $((\,)\ L \twoheadrightarrow_{\mathbb{S}})$, or in other words, the set of $\twoheadrightarrow_{\mathbb{S}}$-reducts of $(\,)\ L$:

$$\begin{aligned}\mathrm{ST}_{\mathbb{S}} \ :\ \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}}) \ &\to\ \wp(\mathrm{Ter}(\boldsymbol{(\lambda_{\mathsf{letrec}})}))\\ L \ &\mapsto\ O \quad \text{where } \langle O, \Phi, \mathrm{src}, \mathrm{tgt}\rangle = ((\,)\ L \twoheadrightarrow_{\mathbb{S}})\end{aligned}$$

§ 1.5.16. The following lemma states that every scope/scope^{+}-delimiting strategy for $Reg^{+}_{\mathsf{letrec}}$ (on $\lambda_{\mathsf{letrec}}$-terms), when restricted to the reducts of a $\lambda_{\mathsf{letrec}}$-term L that expresses a λ-term M, projects to the restriction of a scope^{+}-delimiting strategy for Reg^{+} (on λ-terms) to reducts of M. And it

asserts a similar statement for scope-delimiting strategies. The proof uses the commutation properties described in lemma 1.5.7 between the infinite unfolding $\to_{\triangledown}^{!\omega}$ and the decomposition rewrite relations $\to_{\lambda}$, $\to_{@_0}$, $\to_{@_1}$, $\to_{\mathsf{s}}$, and $\to_{\mathrm{del}}$.

Lemma 1.5.17 (projection of scope-delimiting strategies)**.** Let $\mathbb{S}$ be a scope/scope$^+$-delimiting strategy $\mathbb{S}$ for $\mathit{Reg}_{\mathsf{letrec}}/\mathit{Reg}^+_{\mathsf{letrec}}$, and let L be a $\lambda_{\mathsf{letrec}}$-term that is $\mathbb{S}$-productive. Then there exists a (history-aware) scope/scope$^+$-delimiting strategy $\check{\mathbb{S}}$ for $\mathit{Reg}/\mathit{Reg}^+$ such that the generated sub-ARS $(\llbracket L \rrbracket_{\boldsymbol{\lambda}} \twoheadrightarrow_{\check{\mathbb{S}}})$ of $\llbracket L \rrbracket_{\boldsymbol{\lambda}}$ is the projection (under the unfolding mapping $\llbracket \cdot \rrbracket_{\boldsymbol{\lambda}}$) of the generated sub-ARS $(L \twoheadrightarrow_{\mathbb{S}})$ of L, in the sense that for all L' in $(L \twoheadrightarrow_{\mathbb{S}})$ it holds:

$$L' \twoheadrightarrow_{\mathbb{S}.\triangledown} \cdot \to_{\mathbb{S}.\Delta} L'' \implies \llbracket L' \rrbracket_{\boldsymbol{\lambda}} \to_{\check{\mathbb{S}}} \llbracket L'' \rrbracket_{\boldsymbol{\lambda}}$$
$$\llbracket L' \rrbracket_{\boldsymbol{\lambda}} \to_{\check{\mathbb{S}}} M'' \implies (\exists L'')\ L' \twoheadrightarrow_{\mathbb{S}.\triangledown} \cdot \to_{\mathbb{S}.\Delta} L'' \wedge M'' = \llbracket L'' \rrbracket_{\boldsymbol{\lambda}}$$

As a consequence, $\llbracket L \rrbracket_{\boldsymbol{\lambda}}$ is $\check{\mathbb{S}}$-regular if L is $\mathbb{S}$-regular.

Proof sketch. We can utilise lemma 1.5.7 to make commuting diagrams out of the two formulas above (for any given L'), which allows us to determine $\check{\mathbb{S}}$ with respect to all terms in $(L \twoheadrightarrow_{\mathbb{S}})$. This freedom in the definition of $\check{\mathbb{S}}$ also guarantees the property in the second implication in the lemma. □

Definition 1.5.18 ($\boldsymbol{\mathit{Parse}}^+_{\triangledown}$)**.** By $\boldsymbol{\mathit{Parse}}^+_{\triangledown}$ we denote the CRS comprising the rules of $\boldsymbol{\mathit{Parse}}^+$ as well as the unfolding rules from $\boldsymbol{R}_{\boldsymbol{\triangledown}}$.

Example 1.5.19. When applied to $\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f$, the rewrite relation in $\boldsymbol{\mathit{Parse}}^+_{\triangledown}$ unfolds and decomposes, but at the same time recreates the corresponding λ-term (see also example 1.2.4 and example 1.4.55). In the rewriting sequence below, we prefix the rewrite relations that are due to rules from $\boldsymbol{R}_{\boldsymbol{\triangledown}}$ with $\triangledown$, while the rewrite relations due to rules from $\boldsymbol{\mathit{Parse}}^+_{\triangledown}$ are not

prefixed.

$$
\begin{array}{ll}
 & \mathsf{parse}_0^+(()\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f) \\
\to_{\bigtriangledown.\mathrm{rec}} & \mathsf{parse}_0^+(()\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ \lambda xy.\ f\ y\ x) \\
\to_{\bigtriangledown.\lambda} & \mathsf{parse}_0^+(()\ \lambda x.\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ \lambda y.\ f\ y\ x) \\
\to_{\lambda} & \lambda x.\ \mathsf{parse}_1^+(x, (\lambda x)\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ \lambda y.\ f\ y\ x) \\
\to_{\bigtriangledown.\lambda} & \lambda x.\ \mathsf{parse}_1^+(x, (\lambda x)\ \lambda y.\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f\ y\ x) \\
\to_{\lambda} & \lambda xy.\ \mathsf{parse}_2^+(x, y, (\lambda xy)\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f\ y\ x) \\
\to_{\bigtriangledown.@} & \lambda xy.\ \mathsf{parse}_2^+\left(x, y, (\lambda xy)\ \begin{array}{c}((\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f\ y)) \\ ((\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ x))\end{array}\right) \\
\to_{\bigtriangledown.\mathrm{red}} & \lambda xy.\ \mathsf{parse}_2^+(x, y, (\lambda xy)\ (\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f\ y)\ (\mathsf{let}\ \ \mathsf{in}\ x)) \\
\to_{\bigtriangledown.\mathrm{nil}} & \lambda xy.\ \mathsf{parse}_2^+(x, y, (\lambda xy)\ (\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f\ y)\ x) \\
\to_{@} & \lambda xy.\ \begin{array}{l}((\mathsf{parse}_2^+(x, y, (\lambda xy)\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f\ y))) \\ ((\mathsf{parse}_2^+(x, y, (\lambda xy)\ x)))\end{array} \\
\to_{\mathsf{S}} & \lambda xy.\ \begin{array}{l}((\mathsf{parse}_2^+(x, y, (\lambda xy)\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f\ y))) \\ ((\mathsf{parse}_1^+(x, (\lambda x)\ x)))\end{array} \\
\to_{\mathsf{0}} & \lambda xy.\ (\mathsf{parse}_2^+(x, y, (\lambda xy)\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f\ y))\ x \\
\to_{\bigtriangledown.@} & \lambda xy.\ (\mathsf{parse}_2^+\left(x, y, (\lambda xy)\ \begin{array}{c}((\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f)) \\ ((\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ y))\end{array}\right))\ x \\
\to_{\bigtriangledown.\mathrm{red}} & \lambda xy.\ (\mathsf{parse}_2^+(x, y, (\lambda xy)\ (\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f)\ (\mathsf{let}\ \ \mathsf{in}\ y)))\ x \\
\to_{\bigtriangledown.\mathrm{nil}} & \lambda xy.\ (\mathsf{parse}_2^+(x, y, (\lambda xy)\ (\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f)\ y))\ x \\
\to_{@} & \lambda xy.\ \begin{array}{l}((\mathsf{parse}_2^+(x, y, (\lambda xy)\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f))) \\ ((\mathsf{parse}_2^+(x, y, (\lambda xy)\ y)))\end{array}\ x \\
\to_{\mathsf{0}} & \lambda xy.\ (\mathsf{parse}_2^+(x, y, (\lambda xy)\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f))\ y\ x \\
\to_{\mathsf{S}} & \lambda xy.\ (\mathsf{parse}_1^+(x, (\lambda x)\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f))\ y\ x \\
\to_{\mathsf{S}} & \lambda xy.\ (\mathsf{parse}_0^+(()\ \mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f))\ y\ x \\
\to_{\bigtriangledown.\mathrm{rec}} & \lambda xy.\ \ \ldots\ \ y\ x
\end{array}
$$

Lemma 1.5.20. For all closed $L \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$ the following statements are equivalent:

(i) L expresses an infinite λ-term M, that is, $L \to_{\bigtriangledown}^{\omega} M$.

(ii) $\mathsf{parse}_0^+(()\ L) \to^\omega_{\mathrm{parse}^+_\triangledown} M$, for some infinite λ-term M.

(iii) L is $\mathbb{S}^+$-productive for some scope$^+$-delimiting strategy $\mathbb{S}^+$.

(iv) L is $\mathbb{S}^+$-productive for every scope$^+$-delimiting strategy $\mathbb{S}^+$.

Proof. Let $L \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$. We show the lemma by establishing the implications in the following order: "(iv) $\Rightarrow$ (iii) $\Rightarrow$ (ii) $\Rightarrow$ (i) $\Rightarrow$ (iv)".

The implication "(iv) $\Rightarrow$ (iii)" is clear: (iv) implies that L is productive for e.g. the lazy-unfolding, eager scope$^+$-delimiting strategy for Reg^+_{letrec}.

For showing the implication "(iii) $\Rightarrow$ (ii)", let $\mathbb{S}$ be a scope$^+$-delimiting strategy for Reg^+ such that L is $\mathbb{S}$-productive. Then the strategy $\mathbb{S}$ defines a $\to_{\mathrm{parse}^+_\triangledown}$-rewrite-sequence τ on $\mathsf{parse}_0^+(()\ L)$ by using $\mathbb{S}$ to define next steps on subexpressions that are of the form $\mathsf{parse}_n^+(\ldots,\mathsf{pre}_n([x_1,\ldots,x_n]P))$ in already obtained reducts: if on a term $(\lambda x_1,\ldots,x_n)\ P$ the strategy $\mathbb{S}$ prescribes a $\to_\triangledown$-step, then this step is adopted in τ; if $\mathbb{S}$ prescribes a $\to_\lambda$-step, then τ can continue with a $\to_{\mathrm{parse}^+.\lambda}$-step; if $\mathbb{S}$ prescribes a $\to_{@_0}$- and a $\to_{@_1}$-step, then τ can continue with a $\to_{\mathrm{parse}^+.@}$-step. For the construction of τ, possible steps in subexpressions $\mathsf{parse}_n^+(\ldots,\mathsf{pre}_n([x_1,\ldots,x_n]P))$ at parallel positions have to be interleaved to ensure that the reduction work is done in an outermost-fair way. Productivity of $\mathbb{S}$ on L ensures that always after finitely many steps inside a subexpression $\mathsf{parse}_n^+(\ldots,\mathsf{pre}_n([x_1,\ldots,x_n]P))$ the function symbol parse_n^+ disappears at this position (either entirely, or it is moved deeper over a λ-abstraction or an application). In the terms of the rewrite sequence τ larger and larger λ-term contexts appear at the head. Hence τ is strongly convergent, and it obtains, in the limit, an infinite λ-term; thus it witnesses $\tau : \mathsf{parse}_0^+(()\ L) \to^\omega_{\mathrm{parse}^+_\triangledown} M$.

For the implication "(ii) $\Rightarrow$ (i)", suppose that τ is a rewrite sequence that witnesses $\mathsf{parse}_0^+(()\ L) \to^\omega_{\mathrm{parse}^+_\triangledown} M$ for some infinite λ-term M. Since the $\to_{\mathrm{parse}^+}$-steps require already unfolded parts of the term, they have to 'shadow' unfolding steps. All $\to_\triangledown$-steps in τ take place beneath symbols parse_n^+. So the possibility of $\to_{\mathrm{parse}^+}$-steps during τ depends on the unfolding steps during τ, but not vice versa. Hence a rewrite sequence τ' on $\mathsf{parse}_0^+(()\ L)$ can be constructed that only adopts the $\to_\triangledown$-steps from τ. Since τ is strongly convergent and converges to M, τ' witnesses $\mathsf{parse}_0^+(()\ L) \to^\omega_\triangledown \mathsf{parse}_0^+(()\ M)$. By dropping the 'non-participant' prefix context $\mathsf{parse}_0^+(()\ \Box)$ from all terms in τ', and adapting the steps accordingly, a rewrite sequence τ'' is obtained that witnesses $\tau'' : L \to^\omega_\triangledown M$.

We show the implication "(i) $\Rightarrow$ (iv)" indirectly. So we assume that there is a scope$^+$-delimiting strategy $\mathbb{S}$ such that L is not $\mathbb{S}$-productive. As in the proof

above of "(iii) ⇒ (ii)", $\mathbb{S}$ defines an outermost-fair $\to_{\mathrm{parse}^+_\triangledown}$-rewrite-sequence τ on $\mathsf{parse}^+_0(()\ L)$. But since $\mathbb{S}$ here is a strategy that is not productive for L, it follows that, due to its construction, τ does not succeed in 'pushing' all function symbols letrec to deeper and deeper depth, and thereby building up an infinite λ-term. Instead, this outermost-fair $\to_{\mathrm{parse}^+_\triangledown}$-rewrite-sequence contains infinitely many steps at the position of an outermost occurrence of letrec. Since, other than the $\to_\triangledown$-steps, the $\to_{\mathrm{parse}^+}$-steps (which always take place above outermost occurrences of letrec-symbols) cannot be the reason for this, the same stagnation of an outermost-fair unfolding process takes place if the $\to_{\mathrm{parse}^+}$-steps are postponed, that is dropped from τ. In this way, by again dropping the 'non-participant' prefix context $\mathsf{parse}^+_0(()\ \Box)$ from the terms of τ, and adapting the steps accordingly, we obtain an outermost-fair $\to_\triangledown$-rewrite-sequence starting on $()\ L$ that does not converge to an infinite λ-term. But then lemma 0.7.13 implies that L does not unfold to an infinite λ-term. □

1.6 Proving regularity and strong regularity

§ 1.6.1 (overview)**.** In this section we introduce proof systems that are sound and complete for the notions of regular, and strongly regular λ-terms. In order to prove soundness and completeness, we establish, as auxiliary results, a correspondence between scope/scope$^+$-delimiting strategies for *Reg*/*Reg*$^+$ and closed derivations in the corresponding proof systems. Then we introduce a proof system that is sound and complete for equality between strongly regular λ-terms. Furthermore, we give two proof systems that are sound and complete for the property of $\lambda_{\mathsf{letrec}}$-terms to unfold to λ-terms. And finally, we show the following part of our characterisation result: λ-terms that are unfoldings of $\lambda_{\mathsf{letrec}}$-terms are strongly regular.

We start with a more formal definition of λ-terms and $\lambda_{\mathsf{letrec}}$-terms than definition 1.4.21, by means of derivability in a proof system that formalises term decomposition.

Definition 1.6.2 (λ-terms)**.** We define the set of prefixed λ-terms as those terms in $\mathrm{Ter}(\Sigma^{(\lambda)})$ for which there exists a possibly infinite, completed (see definition 1.3.24) derivation in the proof system $\mathrm{Ter}^\infty(\boldsymbol{\lambda})$ with axioms and rules as shown in fig. 1.13:

$$\mathrm{Ter}^\infty(\boldsymbol{(\lambda)}) := \left\{ M \in \mathrm{Ter}(\Sigma^{(\lambda)}) \;\middle|\; {}^\infty\vdash_{\mathrm{Ter}^\infty(\boldsymbol{\lambda})} M \right\}$$

$$\frac{}{(\lambda\vec{x}y)\ y}\ \mathsf{0} \qquad \frac{(\lambda\vec{x}y)\ M_0}{(\lambda\vec{x})\ \lambda y.\, M_0}\ \lambda \qquad \frac{(\lambda\vec{x})\ M_0 \qquad (\lambda\vec{x})\ M_1}{(\lambda\vec{x})\ M_0\ M_1}\ @$$

$$\frac{(\lambda x_1 \dots x_{n-1})\ M}{(\lambda x_1 \dots x_n)\ M}\ \mathsf{S} \quad \text{(if the binding } \lambda x_n \text{ is vacuous)}$$

Figure 1.13. Proof system $\mathrm{Ter}^{\infty}(\boldsymbol{\lambda})$ for defining the set of λ-terms.

$$\frac{(\lambda\vec{x}f_1 \dots f_n)\ M_0 \qquad \dots \qquad (\lambda\vec{x}f_1 \dots f_n)\ M_n}{(\lambda\vec{x})\ \mathsf{let}\ f_1 = M_1, \dots, f_n = M_n\ \mathsf{in}\ M_0}\ \mathsf{letrec}$$

Figure 1.14. Proof system $\mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$ for defining the set of $\lambda_{\mathsf{letrec}}$-terms defined as an extension of $\mathrm{Ter}^{\infty}(\boldsymbol{\lambda})$ in fig. 1.13 by an additional rule (letrec)

The set of plain λ-terms are those terms that comply with the previous definition when equipped with an empty prefix:

$$\mathrm{Ter}^{\infty}(\boldsymbol{\lambda}) \coloneqq \left\{ M \in \mathrm{Ter}\left(\Sigma^{\lambda}\right) \mid ()\ M \in \mathrm{Ter}^{\infty}(\boldsymbol{(\lambda)}) \right\}$$

Definition 1.6.3 ($\lambda_{\mathsf{letrec}}$-terms)**.** The set of prefixed $\lambda_{\mathsf{letrec}}$-terms comprises those terms out of $\mathrm{Ter}(\Sigma^{(\lambda)}_{\mathsf{letrec}})$ for which there exists a finite derivation in the proof system $\mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$ (fig. 1.14):

$$\mathrm{Ter}(\boldsymbol{(\lambda}_{\mathsf{letrec}}\boldsymbol{)}) \coloneqq \left\{ M \in \mathrm{Ter}(\Sigma^{(\lambda)}_{\mathsf{letrec}}) \mid \vdash_{\mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})} M \right\}$$

The set of plain $\lambda_{\mathsf{letrec}}$-terms are those terms that comply with the previous definition when equipped with an empty prefix:

$$\mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}}) \coloneqq \left\{ M \in \mathrm{Ter}\left(\Sigma^{\lambda}_{\mathsf{letrec}}\right) \mid ()\ M \in \mathrm{Ter}(\boldsymbol{(\lambda}_{\mathsf{letrec}}\boldsymbol{)}) \right\}$$

Building on rules already used in the proof systems for term formation in $\boldsymbol{\lambda}$ and $\boldsymbol{(\lambda)}$ from the definition above, we now introduce proof systems for regularity and strong regularity of λ-terms in $\boldsymbol{\lambda}$.

Definition 1.6.4 (proof systems **Reg**, and $\mathbf{Reg}^{+}$, $\mathbf{Reg}^{+}_{\mathbf{0}}$)**.** The natural-deduction style proof system $\mathbf{Reg}^{+}$ for recognising strongly regular λ-terms contains the axioms and rules as shown in fig. 1.15. In particular, the rule FIX

is a natural-deduction style derivation rule in which marked assumptions from the top of the proof tree can be discharged. Instances of this rule carry the side-condition that the depth $|\mathcal{D}_0|$ of the immediate subderivation $\mathcal{D}_0$ of its premise is greater or equal to 1 (hence this subderivation contains at least one rule instance, and, importantly, for a topmost occurrence of FIX, $\mathcal{D}_0$ must have a bottommost instance of one of the rules (λ), (@), or (S)).

The variant $\mathbf{Reg_0^+}$ of **Reg** contains the same axioms and rules as $\mathbf{Reg^+}$, but in it instances of FIX are subject to the additional side-condition: for all $(\lambda\vec{y})\ N$ on threads in $\mathcal{D}_0$ from open marked assumptions $((\lambda\vec{x})\ M)^u$ downwards it holds that $|\vec{y}| \geq |\vec{x}|$.

The natural-deduction style proof system **Reg** for recognising regular λ-terms differs from $\mathbf{Reg^+}$ by the absence of the rule (S), and the presence instead of the rule (del) in fig. 1.16, and by the restriction of the axiom scheme $(\mathsf{0})$ to the more restricted version displayed in fig. 1.16.

Provability of a term in $\boldsymbol{(\lambda)}$ in one of these proof systems is defined as the existence of a *closed* derivation: for $\boldsymbol{R} \in \{\mathbf{Reg}, \mathbf{Reg^+}, \mathbf{Reg_0^+}\}$ we denote by $\vdash_{\boldsymbol{R}} (\lambda\vec{x})\ M$ the existence of a proof tree $\mathcal{D}$ with conclusion M and with rule instances of $\boldsymbol{R}$ such that all marked assumptions at the top of the $\mathcal{D}$ are discharged at some instance of the rule FIX.

Remark 1.6.5 ($\mathbf{Reg^+}$ versus $\mathbf{Reg_0^+}$)**.** While it will be established in proposition 1.6.14 that provability in $\mathbf{Reg^+}$ and $\mathbf{Reg_0^+}$ coincides, the difference between these systems will come to the fore in annotated versions that are purpose-built for the extraction of $\lambda_{\mathsf{letrec}}$-terms that express λ-terms. This will be explained and illustrated later in example 1.8.6, using annotated versions of the two derivations in example 1.6.8 above.

Remark 1.6.6. The proof system $\mathbf{Reg^+}$ is related to a proof system for nameless, finite terms in the λ-calculus that is used in [47, sec. 2] as part of a translation of λ-terms into 'Lambdascope' interaction nets, which are used for optimal evaluation (in the sense of Lévy) of λ-terms.

The proposition below explains that the side-condition on instances of FIX from the proof systems above to have immediate subderivations $\mathcal{D}_0$ with $|\mathcal{D}_0| \geq 1$ entails a 'guardedness' property for threads from such instances upwards to discharged instances.

Proposition 1.6.7 (cycles are guarded)**.** Let $\mathcal{D}$ be a derivation in **Reg**, in $\mathbf{Reg^+}$ or in $\mathbf{Reg_0^+}$ possibly with open marked assumptions. Then for all instances ι of

$$\frac{(\lambda\vec{x}y)\ M_0}{(\lambda\vec{x})\ \lambda y.\, M_0}\ \lambda \qquad \frac{(\lambda\vec{x})\ M_0 \qquad (\lambda\vec{x})\ M_1}{(\lambda\vec{x})\ M_0\ M_1}\ @$$

$$\frac{}{(\lambda\vec{x}y)\ y}\ \mathsf{0} \qquad \frac{(\lambda x_1 \ldots x_{n-1})\ M}{(\lambda x_1 \ldots x_n)\ M}\ \mathsf{S}\ \begin{array}{l}\text{(if the binding}\\ \lambda x_n \text{ is vacuous)}\end{array}$$

$$\frac{\begin{array}{c}[(\lambda\vec{x})\ M]^u\\ \mathcal{D}_0\\ (\lambda\vec{x})\ M\end{array}}{(\lambda\vec{x})\ M}\ \mathrm{FIX}, u \quad (\text{if } |\mathcal{D}_0| \geq 1)$$

Figure 1.15. The natural-deduction style proof system $\mathbf{Reg}^{+}$ for strongly regular λ-terms is an extension of $\mathrm{Ter}^{\infty}(\boldsymbol{\lambda})$ by one additional rule FIX. In the variant system $\mathbf{Reg}^{+}_{\mathbf{0}}$, instances of FIX are subject to the following side-condition: for all $(\lambda\vec{y})\ N$ on threads in $\mathcal{D}_0$ from open marked assumptions $((\lambda\vec{x})\ M)^u$ downwards it holds that $|\vec{y}| \geq |\vec{x}|$.

$$\frac{(\lambda\vec{x}y)\ M_0}{(\lambda\vec{x})\ \lambda y.\, M_0}\ \lambda \qquad \frac{(\lambda\vec{x})\ M_0 \qquad (\lambda\vec{x})\ M_1}{(\lambda\vec{x})\ M_0\ M_1}\ @$$

$$\frac{}{(\lambda y)\ y}\ \mathsf{0} \qquad \frac{(\lambda x_1 \ldots x_{i-1} x_{i+1} \ldots x_n)\ M}{(\lambda x_1 \ldots x_n)\ M}\ \mathrm{del}\ \begin{array}{l}\text{(if the binding}\\ \lambda x_i \text{ is vacuous)}\end{array}$$

$$\frac{\begin{array}{c}[(\lambda\vec{x})\ M]^u\\ \mathcal{D}_0\\ (\lambda\vec{x})\ M\end{array}}{(\lambda\vec{x})\ M}\ \mathrm{FIX}, u \quad (\text{if } |\mathcal{D}_0| \geq 1)$$

Figure 1.16. The natural-deduction style proof system **Reg** for regular λ-terms arises from the proof system $\mathbf{Reg}^{+}$ by replacing the rule (S) with the rule (del) for the introduction of vacuous bindings in the λ-abstraction prefixes, and by replacing the axiom scheme $(\mathsf{0})$ of $\mathbf{Reg}^{+}$ by the more restricted version here.

the rule FIX in $\mathcal{D}$ it holds: every thread from ι upwards to a marked assumption that is discharged at ι passes at least one instance of a rule (λ) or (@).

Proof. Since for $\mathbf{Reg}^{+}$ and $\mathbf{Reg}_{\mathbf{0}}^{+}$ the argument is analogous, we only consider derivations in **Reg**. So, let $\mathcal{D}$ be a derivation in **Reg**. Furthermore, let ι be an instance of the rule FIX in $\mathcal{D}$ with conclusion $(\lambda\vec{x})\ M$, and let π be a thread from the conclusion of ι upwards to a marked assumption $((\lambda\vec{x})\ M)^{u}$. Let κ be the topmost instance of FIX in $\mathcal{D}$ that is passed on π. By its side-condition, the immediate subderivation of κ has depth greater or equal to 1, and hence there is at least one instance of a rule (λ), (@), or (del) passed on π above κ. If there is an instance of (λ) or (@) on this part of π, we are done. Otherwise only rules (del) are passed on π above κ. But since the rule (del) decreases the length of the abstraction prefix in the term occurrences in a pass from the conclusion to the premise, and since the length of the abstraction prefix at the start of π is the same as at the end of π, namely $|\vec{x}|$, it follows that at least one occurrence of a rule that increases the length of the abstraction prefix must also have been passed on π, on the segment from ι to κ. Since the only rule of **Reg** that increases the length of an abstraction prefix in a pass from conclusion to a premise is the rule (λ), we have also in this case found a desired rule instance on π. □

Example 1.6.8 (difference between $\mathbf{Reg}^{+}$ and $\mathbf{Reg}_{\mathbf{0}}^{+}$)**.** Let M be the infinite unfolding of $\mathsf{let}\ f = \lambda xy.\ f\ y\ x\ \mathsf{in}\ f$ for which we use as a finite representation the equation $M = \lambda xy.\ M\ y\ x$. This term admits the following two derivations in $\mathbf{Reg}^{+}$, with the latter having some redundancy:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{(()\ M)^{u}}{(\lambda x)\ M}\,\mathsf{S}}{(\lambda xy)\ M}\,\mathsf{S} \quad \dfrac{}{(\lambda xy)\ y}\,\mathsf{0}}{(\lambda xy)\ M\ y}\,@ \quad \dfrac{}{(\lambda x)\ x}\,\mathsf{0}}{(\lambda xy)\ x}\,\mathsf{S}}{(\lambda xy)\ M\ y\ x}\,@}{(\lambda x)\ \lambda y.\ M\ y\ x}\,\lambda}{()\ \lambda xy.\ M\ y\ x}\,\lambda}{()\ M}\,\mathrm{FIX}, u$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{((\lambda x)\ \lambda y.\, M\ y\ x)^u}{()\ M}\lambda}{(\lambda x)\ M}\mathsf{S}}{(\lambda xy)\ M}\mathsf{S} \qquad \dfrac{}{(\lambda xy)\ y}\mathsf{0}}{(\lambda xy)\ M\ y}@ \qquad \dfrac{\dfrac{}{(\lambda x)\ x}\mathsf{0}}{(\lambda xy)\ x}\mathsf{S}}{(\lambda xy)\ M\ y\ x}@}{(\lambda x)\ \lambda y.\, M\ y\ x}\lambda}{(\lambda x)\ \lambda y.\, M\ y\ x}\mathrm{FIX}, u}{()\ M}\lambda$$

Note that the first derivation is also a derivation in $\mathbf{Reg}_{\mathbf{0}}^{+}$, but that this is not the case for the secord derivation, as for the occurrence of FIX the side-condition in the system $\mathbf{Reg}_{\mathbf{0}}^{+}$ is violated: on the path from the marked assumption $((\lambda x)\ \lambda y.\, M\ y\ x)^u$ down to the instance of FIX there is the occurrence $()\ M$ of a term with shorter prefix than the term in the assumption and in the conclusion.

See example 1.4.47 for a rewriting sequence in Reg^{+} corresponding to the leftmost path in both derivations, and also fig. 1.12 for the corresponding transition graph.

Example 1.6.9 (difference between ***Reg*** and ***Reg*$^{+}$**)**.** The infinite λ-term from example 1.2.2 with the ***Reg***/***Reg*$^{+}$**-transition-graphs shown in fig. 1.11 is derivable in **Reg** by the following closed derivation using the notation from example 1.4.46:

$$\dfrac{\dfrac{\dfrac{\dfrac{\overbrace{((\lambda b)\ rec_M(b))}^{(\lambda a)\ rec_M(a)})^u}{(\lambda ab)\ rec_M(b)}\mathrm{del} \qquad \dfrac{\dfrac{}{(\lambda a)\ a}\mathsf{0}}{(\lambda ab)\ a}\mathrm{del}}{(\lambda ab)\ rec_M(b)\ a}@}{(\lambda a)\ \lambda b.\, rec_M(b)\ a}\lambda}{(\lambda a)\ rec_M(a)}\mathrm{FIX}, u}{()\ \lambda a.\, rec_M(a)}\lambda$$

When trying to construct a derivation for this term in $\mathbf{Reg}^{+}$ from the bottom upwards, the rules of $\mathbf{Reg}^{+}$ apart from FIX offer only deterministic choices, resulting in an infinite proof tree of the form:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\vdots}{(\lambda abcde)\ rec_M(e)}}{(\lambda abcd)\ \lambda e.\, rec_M(e)}\lambda \quad \dfrac{\dfrac{}{(\lambda abc)\ c}\mathsf{0}}{(\lambda abcd)\ c}\mathsf{S}}{(\lambda abcd)\ rec_M(d)\ c}@}{(\lambda abc)\ \lambda d.\, rec_M(d)\ c}\lambda \quad \dfrac{\dfrac{}{(\lambda ab)\ b}\mathsf{0}}{(\lambda abc)\ b}\mathsf{S}}{(\lambda abc)\ rec_M(c)\ b}@}{(\lambda ab)\ \lambda c.\, rec_M(c)\ b}\lambda \quad \dfrac{\dfrac{}{(\lambda a)\ a}\mathsf{0}}{(\lambda ab)\ a}\mathsf{S}}{(\lambda ab)\ rec_M(b)\ a}@}{(\lambda a)\ \lambda b.\, rec_M(b)\ a}\lambda}{()\ \lambda a.\, rec_M(a)}\lambda$$

But then, since this proof tree does not contain repetitions, the use of the rule FIX in order to discharge assumptions is impossible. Consequently, the term is not derivable in $\mathbf{Reg^+}$.

For the *Reg* and Reg^+ rewriting sequences corresponding to the leftmost paths through the two proofs above, see example 1.4.46. The corresponding transition graphs are displayed in fig. 1.11.

Proposition 1.6.10 (correspondence between proof systems and decomposition CRSs)**.** Let $\mathcal{D}$ be a derivation in $\mathbf{Reg}/\mathbf{Reg^+}$ with conclusion $(\lambda\vec{x})\ M$. Then every path in $\mathcal{D}$ from the conclusion upwards corresponds to a $\to_{\mathrm{reg}}/\to_{\mathrm{reg}^+}$-rewrite-sequence from $(\lambda\vec{x})\ M$: passes over instances of FIX correspond to empty rewrite steps; passes over instances of (@) to the left and to the right correspond to $\to_{@_0}$- and $\to_{@_1}$-steps, respectively; passes over instances of (λ) correspond to $\to_\lambda$-steps; passes over instances of (del)/(S) correspond to $\to_{\mathrm{del}}/\to_{\mathsf{S}}$-steps.

The same holds for (finite or infinite) cyclic paths in $\mathcal{D}$ that return, possibly repeatedly, from a marked assumption at the top down to the conclusion of the instance of FIX at which the respective assumption is discharged.

Proof. The proposition is an easy consequence of the following facts: passes from a term in the conclusion of an instance ι of one of the rules (λ), (del), (S) to the term in the premise of ι correspond to $\to_\lambda$-, $\to_{\mathrm{del}}$-, and $\to_{\mathsf{S}}$-steps, respectively; passes from a term in the conclusion of an instance of (@) to the left and the right premise correspond to $\to_{@_0}$-steps and $\to_{@_1}$-steps, respectively. □

§ 1.6.11 (proofs and scope/scope$^+$-delimiting strategies)**.** Observe that, for the derivation $\mathcal{D}$ in example 1.6.8, the $\to_{\mathrm{reg}^+}$-rewrite-sequences that correspond to paths in $\mathcal{D}$ as described in proposition 1.6.10 are actually rewrite sequences with respect to the *eager* scope$^+$-delimiting strategy $\mathbb{S}^+_{\mathrm{eag}}$ for Reg^+. This illustrates

the general situation, formulated by the lemma below: paths in a derivation $\mathcal{D}$ in $\mathbf{Reg}/\mathbf{Reg}^{+}$ from the conclusion upwards correspond to rewrite sequences according to some – usually history-aware – scope/scope^{+}-delimiting strategy $\mathbb{S}$, which can be extracted from $\mathcal{D}$.

Lemma 1.6.12 (from $\mathbf{Reg}/\mathbf{Reg}^{+}$-derivations to scope/scope^{+}-delimiting strategies)**.** Let $M \in \mathrm{Ter}^{\infty}(\boldsymbol{\lambda})$, and let $\mathcal{D}$ be a closed derivation in $\mathbf{Reg}$ (in $\mathbf{Reg}^{+}$) with conclusion $()\,M$. Then there exists an, in general history-aware, scope-delimiting strategy $\mathbb{S}_{\mathcal{D}}$ for *Reg* (scope^{+}-delimiting strategy $\mathbb{S}_{\mathcal{D}}$ for Reg^{+}) with the following properties:

(i) Every (possibly cyclic) path in $\mathcal{D}$ from the conclusion upwards corresponds to a rewrite sequence with respect to $\mathbb{S}_{\mathcal{D}}$ starting on $()\,M$ in the sense of proposition 1.6.10 where passes over instances of the rules $(@)$ to the left and to the right correspond to $\to_{\mathbb{S}_{\mathcal{D}}.@_0}$-steps and $\to_{\mathbb{S}_{\mathcal{D}}.@_1}$-steps, respectively, and passes over instances of (λ) and of (del) (of (S)) correspond to $\to_{\mathbb{S}_{\mathcal{D}}.\lambda}$-, and $\to_{\mathbb{S}_{\mathcal{D}}.\mathrm{del}}$-steps ($\to_{\mathbb{S}_{\mathcal{D}}.\mathsf{S}}$-steps).

(ii) Every rewrite sequence that starts on $()\,M$ and proceeds according to $\mathbb{S}_{\mathcal{D}}$ corresponds to a (possibly cyclic) path in $\mathcal{D}$ starting at the conclusion in upwards direction: thereby a $\to_{\mathbb{S}_{\mathcal{D}}.@_0}$-step and a $\to_{\mathbb{S}_{\mathcal{D}}.@_1}$-step corresponds to a pass over (possibly successive FIX-instances, or from a marked assumption to the instance of FIX that binds it, followed by) an instance of $(@)$ in direction left and right, respectively; a $\to_{\mathbb{S}_{\mathcal{D}}.\lambda}$-step or $\to_{\mathbb{S}_{\mathcal{D}}.\mathrm{del}}$-step ($\to_{\mathbb{S}_{\mathcal{D}}.\mathsf{S}}$-step) corresponds to a pass over (possibly FIX-instances and assumption bindings to FIX-instances) an instance of (λ) or (del) (of (S)), respectively.

(iii) $\mathrm{ST}_{\mathbb{S}_{\mathcal{D}}}(M) = \{(\lambda\vec{y})\ N \mid \text{the term } (\lambda\vec{y})\ N \text{ occurs in } \mathcal{D}\}$.

Proof. The proof defines a history-aware strategy $\mathbb{S}_{\mathcal{D}}$ for Reg^{+} as a modification of an arbitrary (history-free) strategy for Reg^{+} lifted to a labelled version of Reg^{+}. Thereby the modification is performed according to a given derivation $\mathcal{D}$, and the construction will guarantee that (i), (ii), and (iii) hold.

We establish the lemma only for the case of derivations in $\mathbf{Reg}^{+}$, since the case of derivations in $\mathbf{Reg}$ can be treated analogously. So, let $\mathcal{D}$ be a derivation in $\mathbf{Reg}^{+}$ with conclusion $()\,M$.

In a first step we decorate $\mathcal{D}$ with position labels such that a derivation $\mathcal{D}^{(\mathrm{lb})}$ with conclusion $:(\lambda\epsilon)\,M$ in the variant proof system $\widehat{\mathbf{Reg}^{+}}$ in fig. 1.17 is obtained. Note that the decoration process can be carried out in a bottom-up

$$\frac{l : (\lambda \vec{x} y)\ M_0}{l : (\lambda \vec{x})\ \lambda y.\, M_0}\ \lambda \qquad \frac{l\,0 : (\lambda \vec{x})\ M_0 \qquad l\,1 : (\lambda \vec{x})\ M_1}{l : (\lambda \vec{x})\ M_0\ M_1}\ @$$

$$\frac{}{l : (\lambda \vec{x} y)\ y}\ \mathsf{0} \qquad \frac{l : (\lambda x_1 \ldots x_{n-1})\ M}{l : (\lambda x_1 \ldots x_n)\ M}\ \mathsf{S}\ \text{(if } x_n \text{ does not occur in } M\text{)}$$

$$\frac{\begin{array}{c}[l : (\lambda \vec{x})\ M]^u \\ \mathcal{D}_0 \\ l : (\lambda \vec{x})\ M\end{array}}{l : (\lambda \vec{x})\ M}\ \mathrm{FIX}, u \quad (\text{if } |\mathcal{D}_0| \geq 1)$$

Figure 1.17. Proof system $\widehat{\mathbf{Reg}^+}$ for decorating $\mathbf{Reg}^+$-derivations with labels in $\{0,1\}^*$.

manner, where the label in the conclusion of a rule instance determines the label in the premise(s) if that is not an already labelled term, and where in the case of instances of the rule of FIX also the labels in marked assumptions are determined.

In a second step we use the decorated version $\mathcal{D}^{(\mathrm{lb})}$ of $\mathcal{D}$ to define a history-aware strategy $\mathbb{S}_{\mathcal{D}}$ according to which the term $()\ M$ can be reduced as 'prescribed' by $\mathcal{D}^{(\mathrm{lb})}$. Since the derivations can only determine the strategy $\mathbb{S}_{\mathcal{D}}$ on terms occurring in $\mathcal{D}$, we also have to define $\mathbb{S}_{\mathcal{D}}$ on other terms. This will be done by choosing an arbitrary (but here: history-free) scope$^+$-delimiting strategy $\mathbb{S}$ for *Reg*, and basing the definition of $\mathbb{S}_{\mathcal{D}}$ on it.

We start by defining a labelling of Reg^+ as the ARS for which $\mathbb{S}_{\mathcal{D}}$ will be defined as a history-free strategy, which together with an initial labelling l then yields a history-aware strategy for Reg^+. Assuming $Reg^+ = \langle \mathrm{Ter}^\infty((\boldsymbol{\lambda})), \Phi, \mathrm{src}, \mathrm{tgt}\rangle$ as the formal representation of Reg^+, we define the ARS

$\widehat{Reg^+} := \left\langle \mathrm{Ter}(\mathbf{(\lambda)}), \widehat{\Phi}, \widehat{\mathrm{src}}, \widehat{\mathrm{tgt}} \right\rangle$ where

$$\mathrm{Ter}(\mathbf{(\lambda)}) := \left\{ l : (\lambda \vec{y})\ N \;\middle|\; (\lambda \vec{y})\ N \in \mathrm{Ter}^{\infty}(\mathbf{(\lambda)}), l \in \{0,1\}^* \right\}$$

$$\widehat{\Phi} := \{ \langle l : (\lambda \vec{y})\ N, \phi, l' : (\lambda \vec{y'})\ N' \rangle \mid \text{the following holds} \}$$

there is an instance of (λ), $(@)$, or (S) in $\widehat{\mathbf{Reg^+}}$ with $l : (\lambda \vec{y})\ N$ in the conclusion and the term $l' : (\lambda \vec{y'})\ N'$ in the premise, and with $\phi : (\lambda \vec{y})\ N \to_{\mathrm{reg}^+} (\lambda \vec{y'})\ N'$ (one of) the corresponding step(s) in Reg^+

and where $\widehat{\mathrm{src}}, \widehat{\mathrm{tgt}} : \widehat{\Phi} \to \mathrm{Ter}(\mathbf{(\lambda)})$ are defined as projections on the first, and respectively, the third component of the triples that constitute steps in $\widehat{\Phi}$. Then the relation

$$\begin{aligned} \mathbb{L} := & \left\{ \langle (\lambda y)\ N, l : (\lambda y)\ N \rangle \;\middle|\; (\lambda y)\ N \in \mathrm{Ter}^{\infty}(\mathbf{(\lambda)}), l \in \{0,1\}^* \right\} \\ & \cup \left\{ \langle \phi, \langle (\lambda y)\ N, \phi, (\lambda y)\ N \rangle \rangle \;\middle|\; \langle (\lambda y)\ N, \phi, (\lambda y)\ N \rangle \in \widehat{\Phi} \right\} \end{aligned}$$

is a labelling of Reg^+ to $\widehat{Reg^+}$. As initial labelling we choose the function l that is defined by $l : \mathrm{Ter}^{\infty}(\mathbf{(\lambda)}) \to \mathrm{Ter}(\mathbf{(\lambda)}), (\lambda \vec{x})\ M \mapsto \epsilon : (\lambda \vec{x})\ M.$ and which adds the label 'ϵ'.

Now we define the strategy $\mathbb{S}_{\mathcal{D}}$ with:

$$\mathbb{S}_{\mathcal{D}} := \left\langle \mathrm{Ter}(\mathbf{(\lambda)}), \widehat{\Phi}_{\text{on-}\mathcal{D}^{(\mathrm{lb})}} \cup \widehat{\Phi}_{\text{not-on-}\mathcal{D}^{(\mathrm{lb})}}, \widehat{\mathrm{src}}', \widehat{\mathrm{tgt}}' \right\rangle$$

$$\widehat{\Phi}_{\text{on-}\mathcal{D}^{(\mathrm{lb})}} := \left\{ \langle l : (\lambda \vec{y})\ N, \phi, l' : (\lambda \vec{y'})\ N' \rangle \in \widehat{\Phi} \;\middle|\; \text{the following holds} \right\}$$

there is an instance of (λ), $(@)$, or (S) in $\mathcal{D}^{(\mathrm{lb})}$ with $l : (\lambda \vec{y})\ N$ in the conclusion and the term $l' : (\lambda \vec{y'})\ N'$ in the premise, and with $\phi : (\lambda \vec{y})\ N \to_{\mathrm{reg}^+} (\lambda \vec{y'})\ N'$ (one of) the corresponding step(s) in Reg^+

$$\widehat{\Phi}_{\text{not-on-}\mathcal{D}^{(\mathrm{lb})}} := \left\{ \langle l : (\lambda \vec{y})\ N, \phi, l' : (\lambda \vec{y'})\ N' \rangle \in \widehat{\Phi} \;\middle|\; \text{the following holds} \right\}$$

$l : (\lambda \vec{y})\ N$ does not occur in $\mathcal{D}^{(\mathrm{lb})}$, ϕ is a step according to $\mathbb{S}$

where $\widehat{\mathrm{src}}'$, $\widehat{\mathrm{tgt}}'$ are the appropriate restrictions of $\widehat{\mathrm{src}}$ and $\widehat{\mathrm{tgt}}$.

Note that, by its definition, $\mathbb{S}_{\mathcal{D}}$ is a sub-ARS of $\widehat{Reg^+}$. Now for showing that $\mathbb{S}_{\mathcal{D}}$ is a (history-aware) strategy for $\widehat{Reg^+}$, it has to be established that $\mathbb{S}_{\mathcal{D}}$ is a history-free strategy for the lifted version $\widehat{Reg^+}$ of Reg^+. For this it remains to show that every normal form of $\mathbb{S}_{\mathcal{D}}$ is also a normal form of $\widehat{Reg^+}$. So, let $l : (\lambda \vec{y})\ N \in \mathrm{Ter}^{\infty}(\mathbf{(\lambda)})$ be such that it is not a normal form of $\widehat{Reg^+}$. Then also

$(\lambda\vec{y})\ N$ is not a normal form of Reg^+. For showing that there is a step in $\mathbb{S}_{\mathcal{D}}$ with this labelled term as a source we will distinguish two cases, i.e. whether $l : (\lambda\vec{y})\ N$ occurs on $\mathcal{D}^{(\mathrm{lb})}$ or not.

For the second case, we assume that $l : (\lambda\vec{y})\ N$ does not occur in $\mathcal{D}^{(\mathrm{lb})}$. Then there is a step $\phi : ((\lambda\vec{y})\ N) \to_{\mathbb{S}} ((\lambda\vec{y}')\ N')$ in the scope-delimiting strategy $\mathbb{S}$ in Reg^+, which gives rise to the step $\phi : (l : (\lambda\vec{y})\ N) \to (l' : (\lambda\vec{y}')\ N')$ in $\widehat{Reg^+}$ and in $\mathbb{S}_{\mathcal{D}}$.

For the first case, we assume that $l : (\lambda\vec{y})\ N$ occurs in $\mathcal{D}^{(\mathrm{lb})}$, and we fix an occurrence o. Since by assumption $l : (\lambda\vec{y})\ N$ is not a normal form of $\widehat{Reg^+}$, o cannot be the occurrence of an axiom $(\mathsf{0})$, and hence it is either an occurrence as the conclusion of an instance of one of the rules (λ), $(@)$, (S) in $\mathcal{D}^{(\mathrm{lb})}$, or as a marked assumption in $\mathcal{D}^{(\mathrm{lb})}$. If o is the conclusion of an instance ι of (λ), $(@)$, or (S), then ι defines a step on $l : (\lambda\vec{y})\ N$ which also is a step in $\mathbb{S}_{\mathcal{D}}$. If o is the conclusion of an instance of FIX in $\mathcal{D}$, then we consider an arbitrary path π in $\mathcal{D}^{(\mathrm{lb})}$ from o upwards towards a leaf of $\mathcal{D}^{(\mathrm{lb})}$. Since, due to the side-condition of the rule FIX, immediate subderivations of instances of FIX consist of at least one rule application, π cannot consist merely of applications of FIX. Hence by following π from o upwards, after a number of successive instances of FIX, each of which have $l : (\lambda\vec{y})\ N$ as conclusion and premise, an instance of one of the rules (λ), $(@)$, (S) follows, which witnesses a step with source (λ), $(@)$, (S) in $\widehat{Reg^+}$ and in $\mathbb{S}_{\mathcal{D}}$. Finally, if o is an occurrence in a marked assumption at the top of the proof tree $\mathcal{D}^{(\mathrm{lb})}$, then, since $\mathcal{D}^{(\mathrm{lb})}$ is a closed derivation and due to the form of instances of the assumption-discharging rule FIX, there is also an occurrence o' of $l : (\lambda\vec{y})\ N$ as the conclusion of an instance of FIX in $\mathcal{D}^{(\mathrm{lb})}$. Now the argument above can be applied to the occurrence o' to obtain a step of $\mathbb{S}_{\mathcal{D}}$ on $l : (\lambda\vec{y})\ N$.

By construction $\mathbb{S}_{\mathcal{D}}$ conforms to (i) and (ii) because of the inclusion of $\widehat{\Phi}_{\text{on-}\mathcal{D}^{(\mathrm{lb})}}$ and $\widehat{\Phi}_{\text{not-on-}\mathcal{D}^{(\mathrm{lb})}}$ respectively; (iii) follows from (ii). □

Lemma 1.6.13 (from scope/scope$^+$-delimiting strategies to **Reg**/**Reg**$^+$-derivations)**.** Let $M \in \mathrm{Ter}^{\infty}(\boldsymbol{\lambda})$, and let $\mathbb{S}$ be a scope-delimiting strategy for Reg (a scope$^+$-delimiting strategy for Reg^+) such that $\mathrm{ST}_{\mathbb{S}}(M)$ is finite. Then there exists a closed derivation $\mathcal{D}$ in **Reg** (in $\mathbf{Reg}^{+}_{\mathbf{0}}$, and hence in $\mathbf{Reg}^{+}$) with conclusion $()\ M$ such that the following properties hold (note the minor differences with the items (i), (ii), and (iii) in lemma 1.6.12):

(i) Every (non-cyclic) path in $\mathcal{D}$ from the conclusion upwards to a leaf of the proof tree $\mathcal{D}$ corresponds to a $\to_{\mathbb{S}}$-rewrite-sequence starting on $()\ M$ where passes over instances of the rules $(@)$ to the left and to the right correspond

to $\to_{\mathbb{S}.@_0}$- and $\to_{\mathbb{S}.@_1}$-steps, respectively, and passes over instances of (λ) and of (del) (of (S)) correspond to $\to_{\mathbb{S}.\lambda}$-, and $\to_{\mathbb{S}.\mathrm{del}}$-steps ($\to_{\mathbb{S}.\mathsf{S}}$-steps); passes from the conclusion to the premise of instances of FIX correspond to empty steps.

(ii) Every sufficiently long $\to_{\mathbb{S}}$-rewrite-sequence on $(\lambda\vec{x})\ M$ has an initial segment that corresponds to a (non-cyclic) path in $\mathcal{D}$ from the conclusion upwards to a leaf of the proof tree: thereby a $\to_{\mathbb{S}.@_0}$-step or $\to_{\mathbb{S}.@_1}$-step corresponds to a pass over (possibly some FIX-instances followed by) an instance of $(@)$ in direction left and right, respectively; a $\to_{\mathbb{S}.\lambda}$-step or $\to_{\mathbb{S}.\mathrm{del}}$-step ($\to_{\mathbb{S}.\mathsf{S}}$-step) corresponds to a pass over (possibly some FIX-instances followed by) an instance of (λ) or (del) (of (S)), respectively.

(iii) $\mathrm{ST}_{\mathbb{S}}(M) \subseteq \{(\lambda\vec{y})\ N \mid \text{the term } (\lambda\vec{y})\ N \text{ occurs in } \mathcal{D}\}$.

Proof. We will argue only for the part of the statement of the lemma concerning a scope$^+$-delimiting strategy for *Reg*$^+$, since the case with a scope-delimiting strategy for *Reg* can be established analogously.

Let M be an λ-term, and let $\mathbb{S}$ be a scope$^+$-delimiting strategy for *Reg* such that $\mathrm{ST}_{\mathbb{S}}(M)$ is finite. Now let $\mathcal{D}_0$ be the (trivial) derivation with conclusion $()\ M$, which, in case that this is not an axiom of $\mathbf{Reg_0^+}$ (and $\mathbf{Reg^+}$), is also an assumption, and then is of the form $((()\ M)^u$, carrying an assumption marker u. If $\mathcal{D}_0$ is an axiom, then it is easy to verify that the statements (i), (ii), and (iii) hold.

Otherwise we construct a sequence $\mathcal{D}_1, \mathcal{D}_2, \ldots$ of derivations where each $\mathcal{D}_n$ satisfies the properties (i), (ii), and (iii), where terms in marked assumptions are not also terms in axioms $\mathsf{0}$, and where $\mathcal{D}_{n+1}$ extends $\mathcal{D}_n$ by one additional rule instance above a marked assumption in $\mathcal{D}_n$: For the extension step on a derivation $\mathcal{D}_n$, a marked assumption $((\lambda\vec{y})\ N)^u$ in $\mathcal{D}_n$ is picked with the property that the term $(\lambda\vec{y})\ N$ does not appear in the thread down to the conclusion of $\mathcal{D}_n$.

Suppose that the $\to_{\mathbb{S}}$-rewrite-sequence from the conclusion of $\mathcal{D}_n$ up to the marked assumption is of the form:

$$\tau : ()\ M = (\lambda\vec{x}_0)\ M_0 \to_{\mathbb{S}} (\lambda\vec{x}_1)\ M_1 \to_{\mathbb{S}} \ldots \to_{\mathbb{S}} (\lambda\vec{x}_m)\ M_m = (\lambda\vec{y})\ N$$

Note that, since by assumption $(\lambda\vec{y})\ N$ is not a term in an axiom $\mathsf{0}$ of $\mathbf{Reg_0^+}$, it follows by proposition 1.4.24 (viii), that it is not a $\to_{\mathrm{reg}^+}$-normal-form. Then depending on whether the possible next step(s) in an $\to_{\mathbb{S}}$-rewrite that extends τ by one step is a $\to_{\mathbb{S}.\lambda}$-, $\to_{\mathbb{S}.\mathrm{del}}$-step, or either a $\to_{\mathbb{S}.@_0}$-steps or a $\to_{\mathbb{S}.@_1}$-steps,

the derivation $\mathcal{D}_n$ is extended above the marked assumption $((\lambda\vec{y})\ N)^u$ by an application of λ, S, or @, respectively. For example in the case that τ extends by one additional step to either of the two rewrite sequences:

$$\begin{aligned}\tau_i : ()\ M = (\lambda\vec{x}_0)\ M_0 \to_{\mathbb{S}} \ldots \to_{\mathbb{S}} &(\lambda\vec{x}_m)\ M_m \\ = &(\lambda\vec{x}_m)\ M_{m,0}\ M_{m,1} \to_{\mathbb{S}.@_i} (\lambda\vec{x}_m)\ M_{m,i}\end{aligned}$$

with $i \in \{0,1\}$, the derivation $\mathcal{D}_n$ of the form:

$$\begin{array}{c}\langle((\lambda\vec{x}_m)\ M_{m,0}\ M_{m,1})^u\rangle \\ \mathcal{D}_n \\ ()\ M\end{array}$$

is extended to $\mathcal{D}_{n+1}$:

$$\dfrac{\begin{array}{c}((\lambda\vec{x}_m)\ M_{m,0})^{u_0} \qquad ((\lambda\vec{x}_m)\ M_{m,1})^{u_1} \\ \langle(\lambda\vec{x}_m)\ M_{m,0}\ M_{m,1}\rangle \\ \mathcal{D}_n \\ ()\ M\end{array}}{}\ @$$

for two fresh assumption markers u_0 and u_1 (the angle brackets $\langle\ldots\rangle$ are used here to indicate just a single formula occurrence at the top of the proof tree $\mathcal{D}_n$). If either of $(\lambda\vec{x}_m)\ M_{m,0}$ or $(\lambda\vec{x}_m)\ M_{m,1}$ is an axiom, then the assumption marker is removed and the formula is marked as an axiom $\mathsf{0}$, accordingly. Note that, if the statements (i), (ii), and (iii) are satisfied for $\mathcal{D} = \mathcal{D}_n$, then this is also the case for $\mathcal{D} = \mathcal{D}_{n+1}$. Furthermore, terms in marked assumptions are not terms in axioms of $\mathbf{Reg}_\mathbf{0}^+$.

The extension process continues as long as $\mathcal{D}_n$ contains a marked assumption $((\lambda\vec{y})\ N)^v$ without a '$\mathbf{Reg}_\mathbf{0}^+$-admissible repetition' beneath it, by which we mean the occurrence o of the term $(\lambda\vec{y})\ N$ on the thread down to the conclusion in $\mathcal{D}$, but strictly beneath the marked assumption, such that furthermore all terms on the part of the thread down to o have an abstraction prefix of length greater or equal to $|\vec{y}|$. (Note the connection to the side-condition on instances of the rule FIX in $\mathbf{Reg}_\mathbf{0}^+$, and, in particular, that marked assumptions with an $\mathbf{Reg}_\mathbf{0}^+$-admissible repetition beneath it could be discharged by an appropriately introduced instance of FIX in $\mathbf{Reg}_\mathbf{0}^+$.)

That the extension process terminates can be seen as follows: Suppose that, to the contrary, it continues indefinitely. Then, since the derivation size increases strictly in every step, an infinite proof tree $\mathcal{D}^\infty$ is obtained in the limit, which due to finite branchingness of the proof tree and Kőnig's Lemma

possesses an infinite path π starting at the conclusion. Now note that due to (i), π corresponds to an infinite $\to_{\mathbb{S}}$-rewrite-sequence. Due to proposition 1.4.50, this infinite rewrite sequence must contain a grounded cycle. However, the existence of such grounded cycle contradicts the termination condition of the extension process, because every grounded cycle provides an $\mathbf{Reg}_{\mathbf{0}}^{+}$-admissible repetition.

Let $\mathcal{D}_N$, for some $N \in \mathbb{N}$, be the derivation that is reached when no further extension step, as described, is possible. By the construction the statements (i), (ii), and (iii) are satisfied for $\mathcal{D} = \mathcal{D}_N$. Furthermore, $\mathcal{D}_N$ is a derivation in $\mathbf{Reg}^{+}$ and $\mathbf{Reg}_{\mathbf{0}}^{+}$ in which every leaf at the top is either an axiom $\mathsf{0}$ or an assumption $((\lambda\vec{y})\ N)^u$ marked with a unique marker u, and for every such marked assumption in $\mathcal{D}_N$, there is a $\mathbf{Reg}_{\mathbf{0}}^{+}$-admissible repetition strictly beneath it. This fact enables us to modify $\mathcal{D}_N$ into a closed derivation in $\mathbf{Reg}_{\mathbf{0}}^{+}$ by closing all open assumptions by newly introduced applications of the rule FIX. More precisely, steps of the following kind are carried out repeatedly. A derivation with occurrences of a number of marked assumptions $((\lambda\vec{y})\ N)^{u_i}$ highlighted together with a single occurrence of the term $(\lambda\vec{y})\ N$ in its interior that indicates the $\mathbf{Reg}_{\mathbf{0}}^{+}$-admissible repetition for the displayed marked assumptions:

$$\begin{array}{c} \langle((\lambda\vec{y})\ N)^{u_1}\rangle\ \ldots\ \langle((\lambda\vec{y})\ N)^{u_k}\rangle \\ \mathcal{D}_{000} \\ \langle(\lambda\vec{y})\ N\rangle \\ \mathcal{D}_{00} \\ ()\ M \end{array}$$

is modified into:

$$\begin{array}{c} \langle((\lambda\vec{y})\ N)^{w}\rangle\ \ldots\ \langle((\lambda\vec{y})\ N)^{w}\rangle \\ \mathcal{D}_{000} \\ \dfrac{(\lambda\vec{y})\ N}{\langle(\lambda\vec{y})\ N\rangle}\,\mathrm{FIX}, w \\ \mathcal{D}_{00} \\ ()\ M \end{array}$$

where w is a fresh assumption marker. In every such transformation step the number of open assumptions is strictly decreased, but the properties (i), (ii), and (iii) (for $\mathcal{D}$ the resulting derivation of such a step) is preserved. Hence after finitely many such transformation steps a derivation $\mathcal{D}$ in $\mathbf{Reg}_{\mathbf{0}}^{+}$ without open

assumptions and with the properties (i), (ii), and (iii) is reached, and obtained as the result of this construction. □

As a consequence of the two lemmas above, derivability in $\mathbf{Reg^+}$ and in $\mathbf{Reg_0^+}$ coincides:

Proposition 1.6.14 ($\mathbf{Reg^+} \equiv \mathbf{Reg_0^+}$)**.**

$$\forall M \in \mathrm{Ter}^\infty(\boldsymbol{\lambda}) \quad \vdash_{\mathbf{Reg^+}} () \, M \iff \vdash_{\mathbf{Reg_0^+}} () \, M$$

Proof. The direction "⇐" follows by the fact that every derivation in $\mathbf{Reg_0^+}$ is also a derivation in $\mathbf{Reg^+}$. For the direction "⇒", let M be an infinite term such that $\vdash_{\mathbf{Reg^+}} () \, M$. By lemma 1.6.12 there exists a scope$^+$-delimiting strategy $\mathbb{S}^+$ such that $\mathrm{ST}_{\mathbb{S}^+}(M)$ is finite. But then it follows by lemma 1.6.13 that there is also a closed derivation in $\mathbf{Reg_0^+}$ with conclusion $() \, M$, and hence that $\vdash_{\mathbf{Reg_0^+}} () \, M$. □

Now we have assembled all auxiliary statements that we use for proving a theorem that tightly links derivability in the proof system **Reg** with regularity, and derivability in $\mathbf{Reg^+}$ and in $\mathbf{Reg_0^+}$ with strong regularity, of λ-terms.

Theorem 1.6.15. The following statements hold for the proof systems **Reg**, $\mathbf{Reg^+}$, $\mathbf{Reg_0^+}$:

(i) **Reg** is sound and complete for regular λ-terms. That is, for all $M \in \mathrm{Ter}^\infty(\boldsymbol{\lambda})$ it holds:

$$\vdash_{\mathbf{Reg}} () \, M \qquad \text{if and only if} \qquad M \text{ is regular.}$$

(ii) $\mathbf{Reg^+}$ and $\mathbf{Reg_0^+}$ are sound and complete for strongly regular λ-terms. That is, for all $M \in \mathrm{Ter}^\infty(\boldsymbol{\lambda})$ the following statements are equivalent:

(a) M is strongly regular.

(b) $\vdash_{\mathbf{Reg^+}} () \, M$.

(c) $\vdash_{\mathbf{Reg_0^+}} () \, M$.

Proof. Since the proof of statement (ii) of the theorem can be carried out analogously (taking into account proposition 1.6.14), we argue here only for statement (i).

For "⇒" in (i), let M be a λ-term that is regular. Then there exists a scope-delimiting strategy $\mathbb{S}$ on *Reg* such that $\mathrm{ST}_{\mathbb{S}}(M)$ is finite. By lemma 1.6.13

$$\frac{}{\mathsf{pre}_{n+1}([x_1 \dots x_n y]\, y) = \mathsf{pre}_{n+1}([z_1 \dots z_n u]\, u)}\ 0$$

$$\frac{\mathsf{pre}_n([x_1 \dots x_n]\, s) = \mathsf{pre}_n([z_1 \dots z_n]\, t)}{\mathsf{pre}_{n+1}([x_1 \dots x_n y]\, s) = \mathsf{pre}_{n+1}([z_1 \dots z_n w]\, t)}\ \mathsf{S}$$

(y does not occur free in s; w does not occur free in t)

$$\frac{\mathsf{pre}_{n+1}([x_1 \dots x_n y]\, s) = \mathsf{pre}_{n+1}([z_1 \dots z_n u]\, t)}{\mathsf{pre}_n([x_1 \dots x_n]\, \mathsf{abs}([y]\, s)) = \mathsf{pre}_n([z_1 \dots z_n]\, \mathsf{abs}([u]\, t))}\ \lambda$$

$$\frac{\begin{array}{c}\mathsf{pre}_n([x_1 \dots x_n]\, s_0) = \mathsf{pre}_n([y_1 \dots y_n]\, t_0) \\ \mathsf{pre}_n([x_1 \dots x_n]\, s_1) = \mathsf{pre}_n([y_1 \dots y_n]\, t_1)\end{array}}{\mathsf{pre}_n([x_1 \dots x_n]\, s_0\ s_1) = \mathsf{pre}_n([y_1 \dots y_n]\, t_0\ t_1)}\ @$$

Figure 1.18. Proof system $\mathbf{EQ}^{\infty}_{\alpha}$ for equality of preterms in $\mathbf{(\lambda)}$ modulo $\equiv_\alpha$.

$$\frac{}{(\lambda \vec{x} y)\ y = (\lambda \vec{z} u)\ u}\ 0 \qquad \frac{(\lambda \vec{x} y)\ M = (\lambda \vec{z} u)\ N}{(\lambda \vec{x})\ \lambda y.\, M = (\lambda \vec{z})\ \lambda u.\, N}\ \lambda$$

$$\frac{(\lambda \vec{x})\ M = (\lambda \vec{z})\ N}{(\lambda \vec{x} y)\ M = (\lambda \vec{z} w)\ N}\ \mathsf{S}$$

(if y does not occur in M, and w does not occur in N)

$$\frac{(\lambda \vec{x})\ M_0 = (\lambda \vec{y})\ N_0 \qquad (\lambda \vec{x})\ M_1 = (\lambda \vec{y})\ N_1}{(\lambda \vec{x})\ M_0\ M_1 = (\lambda \vec{y})\ N_0\ N_1}\ @$$

Figure 1.19. Proof system $\mathbf{EQ}^{\infty}$ for equality of terms in $\mathbf{(\lambda)}$ in informal notation.

it follows that there exists a closed derivation $\mathcal{D}$ in **Reg** with conclusion $()\ M$. This derivation witnesses $\vdash_{\mathbf{Reg}} ()\ M$. For "$\Leftarrow$" in (i), suppose that $\vdash_{\mathbf{Reg}} ()\ M$. Then there exists a closed derivation $\mathcal{D}$ in **Reg** with conclusion $()\ M$. Now lemma 1.6.12 entails the existence of a scope-delimiting strategy $\mathbb{S}$ in Reg^{+} such that, in particular, $\mathrm{ST}_{\mathbb{S}}(M)$ is finite. This fact implies that M is regular. □

For defining a proof system for equality of strongly regular λ-terms, we first give a specialised version of the proof system $\mathbf{A}^{\infty}_{\mathcal{K}}$ (for α-equivalence of iCRS

preterms) for **(λ)**-preterms, and a corresponding system on **(λ)**-terms.

Definition 1.6.16 (proof systems $\mathbf{EQ}_{\boldsymbol{\alpha}}^{\boldsymbol{\infty}}$, $\mathbf{EQ}^{\boldsymbol{\infty}}$)**.** The proof system $\mathbf{EQ}_{\boldsymbol{\alpha}}^{\boldsymbol{\infty}}$ for α-equivalence of infinite preterms in **(λ)** consists of the rules displayed in fig. 1.18. The proof system $\mathbf{EQ}^{\boldsymbol{\infty}}$ for equality of infinite terms in **(λ)** consists of the rules displayed in fig. 1.19. Provability in $\mathbf{EQ}_{\boldsymbol{\alpha}}^{\boldsymbol{\infty}}$ and in $\mathbf{EQ}^{\boldsymbol{\infty}}$ is defined, analogous to the proof system $\mathbf{A}_{\mathcal{K}}^{\boldsymbol{\infty}}$ from definition 1.3.24, as the existence of a completed (possibly infinite) derivation, and will, as done for $\mathbf{A}_{\mathcal{S}}^{\boldsymbol{\infty}}$ in definition 1.3.23, be indicated using the symbol ${}^{\infty}\vdash$.

Proposition 1.6.17. The following statements hold for the proof systems $\mathbf{EQ}_{\boldsymbol{\alpha}}^{\boldsymbol{\infty}}$ and $\mathbf{EQ}^{\boldsymbol{\infty}}$.

(i) $\mathbf{EQ}_{\boldsymbol{\alpha}}^{\boldsymbol{\infty}}$ is sound and complete for $\equiv_\alpha$ on **(λ)**-preterms. That is, the following holds for all closed preterms s and t in **(λ)**:

$$ {}^{\infty}\vdash_{\mathbf{EQ}_{\boldsymbol{\alpha}}^{\boldsymbol{\infty}}} \mathsf{pre}_0(s) = \mathsf{pre}_0(t) \qquad \text{if and only if} \qquad s \equiv_\alpha t. $$

(ii) $\mathbf{EQ}^{\boldsymbol{\infty}}$ is sound and complete for equality between **(λ)**-terms. That is, the following holds for all terms M and N in **(λ)**:

$$ {}^{\infty}\vdash_{\mathbf{EQ}^{\boldsymbol{\infty}}} M = N \qquad \text{if and only if} \qquad M = N. $$

Proof. For statement (i) it suffices to show that, for an equation $\mathsf{pre}_0(s) = \mathsf{pre}_0(t)$ between preterms of **(λ)**, derivability in $\mathbf{EQ}_{\boldsymbol{\alpha}}^{\boldsymbol{\infty}}$ coincides with derivability of this equation in the general proof system $\mathbf{A}_{\mathcal{K}}^{\boldsymbol{\infty}}$ for α-equivalence between iCRS preterms in definition 1.3.24. Given a derivation $\mathcal{D}^{\infty}$ in $\mathbf{EQ}_{\boldsymbol{\alpha}}^{\boldsymbol{\infty}}$, a derivation $\mathcal{D}_{\mathcal{K}}^{\infty}$ in $\mathbf{A}_{\mathcal{K}}^{\boldsymbol{\infty}}$ results by replacing each formula occurrence $\mathsf{pre}_n([x_1 \dots x_n]s) = \mathsf{pre}_n([y_1 \dots y_n]t)$ by the formula occurrence $\{x_1 \dots x_n\}s = \{y_1 \dots y_n\}t$ and adding an instance of the rule for the function symbol pre_0 at the bottom. Then instances of the axioms and rules (0), (@), (λ), and (S) in $\mathcal{D}^{\infty}$ correspond to instances of axioms and rules (0), (app), ([ ]), and (S) in $\mathcal{D}_{\mathcal{K}}^{\infty}$, respectively. This proof transformation also has an inverse.

For "⇐" in statement (ii) it suffices to note: Every scope$^+$-delimiting strategy for Reg^+ can be used to stepwise extend finite derivations in $\mathbf{EQ}^{\boldsymbol{\infty}}$ with conclusion $() M = () M$ by one additional rule application above a leaf containing a formula $(\lambda y) N = (\lambda y) N$ that is not an axiom 0, which implies that $(\lambda y) N$ is not a normal form of $\to_{\mathrm{reg}^+}$. If these extensions are carried out in a fair manner by extending all non-axiom leafs at depth n in the proof tree before proceeding with leafs at depth $> n$, then in the limit a completed derivation in $\mathbf{EQ}^{\boldsymbol{\infty}}$ is obtained.

$$\dfrac{\begin{array}{c}[(\lambda\vec{x})\ M = (\lambda\vec{y})\ N]^u \\ \mathcal{D}_0 \\ (\lambda\vec{x})\ M = (\lambda\vec{y})\ N\end{array}}{(\lambda\vec{x})\ M = (\lambda\vec{y})\ N}\ \mathrm{FIX}, u \quad (\text{if } |\mathcal{D}_0| \geq 1)$$

Figure 1.20. The rule FIX, which is added to the rules of $\mathbf{EQ}^{\infty}$ from fig. 1.19 in order to obtain the proof system $\mathbf{Reg}^{+}_{=}$ for equality of strongly regular λ-terms.

For "$\Rightarrow$" in statement (ii), suppose that $\mathcal{D}^{\infty}$ is a completed derivation in $\mathbf{EQ}^{\infty}$ with conclusion $()\ M = ()\ N$. Let $\mathsf{pre}_0(s)$ and $\mathsf{pre}_0(t)$ be preterm representatives of $()\ M$ and $()\ N$, respectively. Now a completed derivation $\mathcal{D}^{\infty}_{\mathrm{pter}}$ in $\mathbf{EQ}^{\infty}_{\boldsymbol{\alpha}}$ can be found by developing it step by step from the conclusion $\mathsf{pre}_0(s) = \mathsf{pre}_0(t)$ upwards, parallel to $\mathcal{D}^{\infty}$, and following the rules of $\mathbf{EQ}^{\infty}_{\boldsymbol{\alpha}}$, which are invertible (that is, the premises of a rule instance are uniquely determined by the conclusion). Then $\mathcal{D}^{\infty}_{\mathrm{pter}}$ is a preterm representative version of $\mathcal{D}^{\infty}$. By using (i), it follows that $s \equiv_{\alpha} t$. Since s and t are preterm representatives of M and N, respectively, $M = N$ follows. □

Definition 1.6.18 (the proof system $\mathbf{Reg}^{+}_{=}$)**.** The natural-deduction-style proof system $\mathbf{Reg}^{+}_{=}$ for equality of strongly regular λ-terms has all the rules of the proof system $\mathbf{EQ}^{\infty}$ from definition 1.6.16 and fig. 1.19, and additionally, the rule FIX in fig. 1.20. But contrary to the definition in $\mathbf{EQ}^{\infty}$, provability of an equation $(\lambda\vec{x})\ M = (\lambda\vec{x})\ N$ in $\mathbf{Reg}^{+}_{=}$ is defined as the existence of a finite closed derivation with conclusion $(\lambda\vec{x})\ M = (\lambda\vec{x})\ N$.

Theorem 1.6.19. $\mathbf{Reg}^{+}_{=}$ is sound and complete for equality between strongly regular λ-terms. That is, for all strongly regular λ-terms M and N it holds:

$$\vdash_{\mathbf{Reg}^{+}_{=}} M = N \qquad \text{if and only if} \qquad M = N.$$

Proof sketch. Let M and N be strongly regular λ-terms. In view of proposition 1.6.17 (ii), it suffices to show:

$$\vdash_{\mathbf{Reg}^{+}_{=}} M = N \qquad \text{if and only if} \qquad {}^{\infty}\vdash_{\mathbf{EQ}^{\infty}} M = N. \tag{1.1}$$

For showing "$\Leftarrow$" in (1.1), let $\mathcal{D}^{\infty}$ be a derivation in $\mathbf{EQ}^{\infty}$ with conclusion $()\ M = ()\ N$. Since paths in $\mathcal{D}^{\infty}$ correspond to $\rightarrow_{\mathrm{reg}^{+}}$-rewrite-sequences,

$$\frac{\{[(\lambda\vec{x})\ \mathsf{c}_{f_i}]^{u_i}\}_{i=1,\ldots,n} \quad \mathcal{D}_j \quad \{\ldots\ (\lambda\vec{x})\ L_j[\vec{f} \coloneqq \vec{\mathsf{c}}_{\vec{f}}]\ \ldots\}_{j=0,\ldots,n}}{(\lambda\vec{x})\ \mathsf{let}\ f_1 = L_1 \ldots f_n = L_n\ \mathsf{in}\ L_0}\ \mathrm{FIX}_{\mathsf{letrec}}, u_1, \ldots, u_n$$

where $\mathsf{c}_{f_1}, \ldots, \mathsf{c}_{f_n}$ are distinct constants fresh for $L_1, \ldots, L_n$, and substitutions $L_j[\vec{f} \coloneqq \vec{\mathsf{c}}_{\vec{f}}]$ stands for $L_j[f_1 \coloneqq \mathsf{c}_{f_1}, \ldots, f_n \coloneqq \mathsf{c}_{f_n}]$.

side-conditions: $|\vec{y}| \geq |\vec{x}|$ holds for the prefix length of every $(\lambda\vec{y})\ N$ on a thread in $\mathcal{D}_j$ for $j \in \{0, \ldots, n\}$ from an open assumptions $((\lambda\vec{x})\ \mathsf{c}_{f_i})^{u_i}$ downwards; for bottommost instances: the arising derivation is guarded on access path cycles.

Figure 1.21. The rule $\mathrm{FIX}_{\mathsf{letrec}}$ for the natural-deduction style proof systems $\mathbf{Reg}_{\mathsf{letrec}}$ and $\mathbf{Reg}^{+}_{\mathsf{letrec}}$ on $\lambda_{\mathsf{letrec}}$-terms.

and since the number of generated subterms of both M and N are finite (as a consequence of their strong regularity), on every infinite thread equation repetitions occur. These repetitions can be used to cut all infinite threads by appropriate introductions of instances of FIX in order to obtain a finite and closed derivation in $\mathbf{Reg}^{+}_{=}$ with the same conclusion.

For showing "$\Rightarrow$" in (1.1), let $\mathcal{D}$ be a closed derivation in $\mathbf{Reg}^{+}_{=}$ with conclusion $()\ M = ()\ N$. Now $\mathcal{D}$ can be unfolded into an infinite derivation $\mathcal{D}^{\infty}$ in $\mathbf{EQ}^{\infty}$ by repeatedly removing a bottommost instance of FIX and inserting its immediate subderivation above each of the marked assumptions the instance discharges. If this process is organised in a fair manner with respect to bottommost instances of FIX, then in the limit an infinite completed proof tree with conclusion $()\ M = ()\ N$ in $\mathbf{EQ}^{\infty}$ is obtained. For productivity of this process it is decisive that the side-condition on every instance ι of the rule FIX guarantees that on threads from the conclusion of ι to a marked assumption that is discharged by ι at least one instance of a rule different from FIX is passed. □

For the purpose of the following definition and the respective proof system in fig. 1.21 we extend the signature $\Sigma^{\lambda}_{\mathsf{letrec}}$ of $\mathbf{Reg}_{\mathsf{letrec}}$ and $\mathbf{Reg}^{+}_{\mathsf{letrec}}$ by an infinite

set of constants for which we use the symbol c as syntactical variables which frequently carry index subscripts.

Definition 1.6.20 (proof systems $\mathbf{Reg}_{\mathsf{letrec}}$, $\mathbf{Reg}^{+}_{\mathsf{letrec}}$, and $\mathbf{ug-Reg}_{\mathsf{letrec}}$, $\mathbf{ug-Reg}^{+}_{\mathsf{letrec}}$)**.** The proof systems $\mathbf{Reg}^{+}_{\mathsf{letrec}}$ and $\mathbf{Reg}_{\mathsf{letrec}}$ for $\lambda_{\mathsf{letrec}}$-terms arise from the proof systems $\mathbf{Reg}^{+}$ and $\mathbf{Reg}$ (see definition 1.6.4, fig. 1.15 and fig. 1.16), respectively, by replacing the terms in the axioms $(\mathsf{0})$, and rules (λ), $(@)$, (S) and (del) through $\lambda_{\mathsf{letrec}}$-terms with abstraction prefixes accordingly, and by replacing the rule FIX with the rule $\mathrm{FIX}_{\mathsf{letrec}}$ in fig. 1.21. The side-condition concerning access path cycles on the derivation arising by an instance of $\mathrm{FIX}_{\mathsf{letrec}}$ pertains only to bottommost occurrences of this rule, and is explained below. By $\mathrm{FIX}^{-}_{\mathsf{letrec}}$ we mean the variant of the rule $\mathrm{FIX}_{\mathsf{letrec}}$ in which the side-condition concerning guardedness of the arising derivation on access path cycles has been dropped. By $\mathbf{ug-Reg}_{\mathsf{letrec}}/\mathbf{ug-Reg}^{+}_{\mathsf{letrec}}$ we denote the variants of $\mathbf{Reg}_{\mathsf{letrec}}/\mathbf{Reg}^{+}_{\mathsf{letrec}}$, respectively, in which the rule $\mathrm{FIX}_{\mathsf{letrec}}$ is replaced by the rule $\mathrm{FIX}^{-}_{\mathsf{letrec}}$.

Let $\mathcal{D}$ be a derivation in one of these proof systems. By an *access path* of $\mathcal{D}$ we mean a (possibly cyclic) path π in $\mathcal{D}$ such that:

(a) π starts at the conclusion and can proceed in upwards direction;

(b) at instances of $(@)$, π can step from the conclusion to one of the premises;

(c) at instances of $\mathrm{FIX}_{\mathsf{letrec}}$, π can step from the conclusion to the rightmost premise (which corresponds to the body of the let-expression);

(d) when arriving at a marked assumption $((\lambda\vec{x})\ \mathsf{c}_{f_i})^{u_i}$ that is discharged at an application of $\mathrm{FIX}_{\mathsf{letrec}}$ of the form as displayed in fig. 1.21, π can step over to the conclusion $(\lambda\vec{x})\ L_i[\vec{f} := \vec{\mathsf{c}}_{\vec{f}}]$ of the subderivation $\mathcal{D}_i$ of that application of $\mathrm{FIX}_{\mathsf{letrec}}$, and proceed from there, again in upwards direction.

For every formula occurrence o in $\mathcal{D}$, by a *relative access path* from o we mean a path with the properties (b)–(d) that starts at o and proceeds in upwards direction. An access path (or relative access path) in $\mathcal{D}$ is *cyclic* if there is a formula occurrence in $\mathcal{D}$ that is visited more than once.

Now we say that $\mathcal{D}$ is *guarded on access path cycles* if every cyclic access path contains, on each of its cycles, at least one *guard*, that is, an instance of a rule (λ) or $(@)$. We say that $\mathcal{D}$ is *guarded* if every relative access path contains a guard on each of its cycles.

Example 1.6.21. The $\lambda_{\mathsf{letrec}}$-term $L_1 = \mathsf{let}\ f = f,\ g = \lambda x.\, x\ \mathsf{in}\ \lambda y.\, g$ admits the following closed derivation $\mathcal{D}_1$ in $\mathbf{Reg}_{\mathsf{letrec}}/\mathbf{Reg}^{+}_{\mathsf{letrec}}$:

$$\dfrac{\dfrac{\dfrac{(()\ \mathsf{c}_g)^{u_2}}{(\lambda y)\ \mathsf{c}_g}\,\mathrm{del}/\mathsf{S}}{()\ \lambda y.\,\mathsf{c}_g}\,\lambda \qquad (()\ \mathsf{c}_f)^{u_1} \qquad \dfrac{\dfrac{}{(\lambda x)\ x}\,0}{()\ \lambda x.\,x}\,\lambda}{()\ \mathsf{let}\ f = f,\ g = \lambda x.\,x\ \mathsf{in}\ \lambda y.\,g}\,\mathrm{FIX}_{\mathsf{letrec}}, u_1, u_2$$

This derivation can be built in a straightforward way, from the bottom upwards. Note that $\mathcal{D}$ is guarded on access path cycles, and hence that the instance of $\mathrm{FIX}_{\mathsf{letrec}}$ at the bottom is a valid one, because: $\mathcal{D}$ does not possess any cyclic access paths. In particular, there is no access path in $\mathcal{D}_1$ that reaches the first premise of the instance of $\mathrm{FIX}_{\mathsf{letrec}}$: this premise is the starting point of an unguarded relative access path, which entails that $\mathcal{D}_1$ itself is not guarded.

Now consider the $\lambda_{\mathsf{letrec}}$-term $L_2 = \mathsf{let}\ f = f,\ g = \lambda x.\,x\ \mathsf{in}\ \lambda y.\,f\ g$. When trying to construct a derivation in $\mathbf{Reg}_{\mathsf{letrec}}/\mathbf{Reg}^{+}_{\mathsf{letrec}}$ for this term in a bottom-up manner, one arrives at the closed derivation $\mathcal{D}_2$ in $\mathbf{ug}-\mathbf{Reg}_{\mathsf{letrec}}/\mathbf{ug}-\mathbf{Reg}^{+}_{\mathsf{letrec}}$:

$$\dfrac{\dfrac{\dfrac{\dfrac{(()\ \mathsf{c}_f)^{u_1}}{(\lambda y)\ \mathsf{c}_f}\,\mathrm{del}/\mathsf{S} \qquad \dfrac{(()\ \mathsf{c}_g)^{u_2}}{(\lambda y)\ \mathsf{c}_g}\,\mathrm{del}/\mathsf{S}}{(\lambda y)\ \mathsf{c}_f\ \mathsf{c}_g}\,@}{()\ \lambda y.\,\mathsf{c}_f\ \mathsf{c}_g}\,\lambda \qquad (()\ \mathsf{c}_f)^{u_1} \qquad \dfrac{\dfrac{}{(\lambda x)\ x}\,0}{()\ \lambda x.\,x}\,\lambda}{()\ \mathsf{let}\ f = f,\ g = \lambda x.\,x\ \mathsf{in}\ \lambda y.\,f\ g}\,\mathrm{FIX}^{-}_{\mathsf{letrec}}, u_1, u_2$$

However, $\mathcal{D}_2$ is not a valid derivation in $\mathbf{Reg}_{\mathsf{letrec}}/\mathbf{Reg}^{+}_{\mathsf{letrec}}$, as the inference step at the bottom is an instance of $\mathrm{FIX}^{-}_{\mathsf{letrec}}$, but not of $\mathrm{FIX}_{\mathsf{letrec}}$, because the side-condition on the arising derivation to be guarded on access path cycles is not satisfied: now there is an access path that reaches the first premise of the derivation and that continues looping on this an unguarded cycle. Since the bottom-up search procedure for derivations is deterministic in this case, it follows that $()\ L_2$ is not derivable in $\mathbf{Reg}$ nor in $\mathbf{Reg}^{+}$.

Similar as the correspondence, stated by proposition 1.6.10, between (possibly cyclic) paths in a derivation in $\mathbf{Reg}$ and $\mathbf{Reg}^{+}$ starting at the conclusion and rewrite sequences with respect to $\to_{\mathrm{reg}}$ and $\to_{\mathrm{reg}^{+}}$ on the infinite term in the conclusion, there is also the following correspondence between access paths in a derivation in $\mathbf{Reg}_{\mathsf{letrec}}$ and $\mathbf{Reg}^{+}_{\mathsf{letrec}}$, and rewrite sequences with respect to $\to_{\mathrm{reg}}$ and $\to_{\mathrm{reg}^{+}}$ on the $\lambda_{\mathsf{letrec}}$-term in the conclusion.

Proposition 1.6.22. Let $\mathcal{D}$ be a derivation in $\mathbf{Reg}_{\mathsf{letrec}}$ or $\mathbf{ug}-\mathbf{Reg}_{\mathsf{letrec}}$ (in $\mathbf{Reg}^{+}_{\mathsf{letrec}}$ or in $\mathbf{ug}-\mathbf{Reg}^{+}_{\mathsf{letrec}}$) with conclusion $(\lambda\vec{x})\ L$.

Then every access path in $\mathcal{D}$ to an occurrence o of a term $(\lambda\vec{y})\ P$ corresponds to a $\to_{\mathrm{reg}}$-rewrite-sequence $(\lambda\vec{x})\ L \twoheadrightarrow_{\mathrm{reg}} (\lambda\vec{y})\ \mathsf{let}\ B\ \mathsf{in}\ \tilde{P}$ (to a $\to_{\mathrm{reg}^+}$-rewrite-sequence $(\lambda\vec{x})\ L \twoheadrightarrow_{\mathrm{reg}^+} (\lambda\vec{y})\ \mathsf{let}\ B\ \mathsf{in}\ \tilde{P}$), where B arises as the union of all outermost binding groups in conclusions of instances of $\mathrm{FIX}_{\mathsf{letrec}}$ below o, and $(\lambda\vec{y})\ P = (\lambda\vec{y})\ \tilde{P}[\vec{f} \coloneqq \vec{\mathsf{c}}_{\vec{f}}]$ where $\vec{f}$ is comprised of the function variables occurring in B and $\vec{\mathsf{c}}_{\vec{f}}$ distinct constants for $\vec{f}$ as chosen by $\mathcal{D}$. More precisely:

(a) a pass over an instance of $\mathrm{FIX}_{\mathsf{letrec}}$ corresponds to an empty or $\to_{\triangledown.\mathrm{letrec}}$-step, dependent on whether the instance is the bottommost $\mathrm{FIX}_{\mathsf{letrec}}$-instance or not;

(b) a pass over an instance of the rule (@) to the left/to the right corresponds to a $\to_{@_0}$-step/$\to_{@_1}$-step, which, if the application is somewhere above an instance of $\mathrm{FIX}_{\mathsf{letrec}}$, has to be preceded by a $\to_{\triangledown.@}$-step;

(c) a pass over an instance of the rule (λ) corresponds to a $\to_{\lambda}$-step which, if the application is above an instance of $\mathrm{FIX}_{\mathsf{letrec}}$, has to be preceded by a $\to_{\triangledown.\lambda}$-step;

(d) a pass over an instance of the rule (del)/(S) corresponds to a $\to_{\mathrm{del}}$/$\to_{\mathsf{S}}$-step, possibly preceded by an application of $\to_{\triangledown.\mathrm{red}}$.

(e) a step from a marked assumption to a premise of a $\mathrm{FIX}_{\mathsf{letrec}}$-instances, a step as described in (c) of the definition of access paths, corresponds to an $\to_{\triangledown.\mathrm{rec}}$-step followed by a $\to_{\triangledown.\mathrm{red}}$-step.

Example 1.6.23. The $\lambda_{\mathsf{letrec}}$-term $\mathsf{let}\ f = \mathsf{let}\ g = f\ \mathsf{in}\ g\ \mathsf{in}\ f$ from example 0.3.10 does not unfold to a λ-term, as can be regognised considering this derivation:

$$\dfrac{\dfrac{(\mathsf{c}_f)^u \qquad \dfrac{(\mathsf{c}_f)^u \qquad (\mathsf{c}_g)^v}{()\ \mathsf{let}\ g = f\ \mathsf{in}\ g}\,\mathrm{FIX}_{\mathsf{letrec}}, v}{()\ \mathsf{let}\ f = \mathsf{let}\ g = f\ \mathsf{in}\ g\ \mathsf{in}\ f}\,\mathrm{FIX}^{-}_{\mathsf{letrec}}, u$$

The instance of $\mathrm{FIX}^{-}_{\mathsf{letrec}}$ at the bottom is not an instance of $\mathrm{FIX}_{\mathsf{letrec}}$, since it is not guarded (has an unguarded cyclic access path that reaches and cycles on the left premise of the instance of $\mathrm{FIX}_{\mathsf{letrec}}$.

Lemma 1.6.24. Let $\mathcal{D}$ be a closed derivation in $\mathbf{ug-Reg}_{\mathsf{letrec}}$ (in $\mathbf{ug-Reg}^{+}_{\mathsf{letrec}}$) with conclusion $()\ L$. Then there exists a scope-delimiting (scope$^+$-delimiting) strategy $\mathbb{S}_{\mathcal{D}}$ for Reg_{letrec} (for $Reg^{+}_{\mathsf{letrec}}$) with the following properties:

(i) Every access path in $\mathcal{D}$ corresponds to a rewrite sequence with respect to $\mathbb{S}_{\mathcal{D}}$ starting on $()\ L$ in the sense of proposition 1.6.22.

(ii) Every rewrite sequence that starts on $()\ L$ and proceeds according to $\mathbb{S}_{\mathcal{D}}$ corresponds to an access path in $\mathcal{D}$ with correspondences as described in proposition 1.6.22.

(iii) $\mathrm{ST}_{\mathbb{S}_{\mathcal{D}}}(L) = \left\{ (\lambda\vec{y})\ \mathsf{let}\ B\ \mathsf{in}\ \tilde{P} \;\middle|\; \begin{array}{l} (\lambda\vec{y})\ \mathsf{let}\ B\ \mathsf{in}\ \tilde{P} \text{ arises from an} \\ \text{occurrence of } (\lambda\vec{y})\ P \text{ on an} \\ \text{access path of } \mathcal{D} \text{ as described} \\ \text{in proposition 1.6.22} \end{array} \right\}$.
As a consequence of that $\mathcal{D}$ is finite, L is $\mathbb{S}_{\mathcal{D}}$-regular.

(iv) L is $\mathbb{S}_{\mathcal{D}}$-productive $\Leftrightarrow$ $\mathcal{D}$ is guarded (i.e. $\mathcal{D}$ derivation in $\mathbf{Reg}_{\mathsf{letrec}}$ ($\mathbf{Reg}^{+}_{\mathsf{letrec}}$)).

Proof. Given a closed derivation $\mathcal{D}$ with conclusion $()\ L$ (for example) in $\mathbf{ug-Reg}^{+}_{\mathsf{letrec}}$, a scope^{+}-delimiting strategy $\mathbb{S}_{\mathcal{D}}$ for $Reg^{+}_{\mathsf{letrec}}$ such that (i)–(iv) hold can be extracted from $\mathcal{D}$ similar as in the proof of a scope^{+}-delimiting strategy $\mathbb{S}_{\mathcal{D}}$ in Reg^{+} was extracted from a closed derivation in $\mathbf{Reg}^{+}$. That the extracted strategy $\mathbb{S}_{\mathcal{D}}$ is productive (not productive) for L if $\mathcal{D}$ is guarded (not guarded) can be seen by the fact that $\mathbb{S}_{\mathcal{D}}$-rewrite-sequences correspond to access paths of $\mathcal{D}$ in the sense as stated by proposition 1.6.22. □

Now we will prove that derivability in $\mathbf{ug-Reg}_{\mathsf{letrec}}/\mathbf{ug-Reg}^{+}_{\mathsf{letrec}}$ is guaranteed for all $\lambda_{\mathsf{letrec}}$-terms, and that derivability in $\mathbf{Reg}_{\mathsf{letrec}}/\mathbf{Reg}^{+}_{\mathsf{letrec}}$ is a property of a $\lambda_{\mathsf{letrec}}$-term that is decidable by an easy parsing process.

Proposition 1.6.25. The following statements hold:

(i) For every $\lambda_{\mathsf{letrec}}$-term L, $()\ L$ is derivable both in $\mathbf{ug-Reg}_{\mathsf{letrec}}$ and in $\mathbf{ug-Reg}^{+}_{\mathsf{letrec}}$.

(ii) For every $\lambda_{\mathsf{letrec}}$-term L, derivability of $()\ L$ in $\mathbf{Reg}^{+}_{\mathsf{letrec}}$ is decidable in at most quadratic time in the size of L.

Proof. For (i) note that for every $\lambda_{\mathsf{letrec}}$-term L, a closed derivation $\mathcal{D}_L$ with conclusion $()\ L$ in $\mathbf{ug-Reg}^{+}_{\mathsf{letrec}}$ can be produced by a bottom-up construction following the term structure of L. Hereby use of the rules (S) can be restricted to instances immediately below marked assumptions such that, viewed from a (non-cyclic) path π from the conclusion upwards to a marked assumption, these (S)-instances are only introduced to shorten the frozen abstraction prefixes by

all λ-abstractions that have become frozen on π (in order to conform to the side-condition on $\mathrm{FIX}^-_{\mathsf{letrec}}$-instances to have the same frozen abstraction prefix lengths in the discharged marked assumptions as in the conclusion and in the premises).

Now for (ii) in order to decide derivability of $()\ L$ in $\mathbf{Reg}^+_{\mathsf{letrec}}$, it suffices to decide whether the derivation $\mathcal{D}_L$ in $\mathbf{ug} - \mathbf{Reg}^+_{\mathsf{letrec}}$ obtained as described above, or its bottommost instance of $\mathrm{FIX}^-_{\mathsf{letrec}}$ if there is any, is guarded on all of its access path cycles. (Note that in the construction of $\mathcal{D}_L$ only the freedom in placing instances of (S) has been used in a certain, namely lazy, way. The specific placement of instances of these rules does not interfere with the existence or non-existence of guards, that is instances of λ or @ on cycles of access paths.) For this it remains to check whether every cycle on an access path in $\mathcal{D}_L$ has a guard. This can be done by exploring the proof tree of $\mathcal{D}_L$ according to all possible access paths (until for the first time a cycle is concluded) and checking for the existence of guards on cycles. □

We now can prove soundness and completeness of the proof system $\mathbf{Reg}^+_{\mathsf{letrec}}$ for the property of $\lambda_{\mathsf{letrec}}$-terms to unfold to λ-terms.

Theorem 1.6.26. $\mathbf{Reg}^+_{\mathsf{letrec}}$ is sound and complete for the property of $\lambda_{\mathsf{letrec}}$-terms to unfold to a λ-term. That is, for every term $L \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$ the following statements are equivalent:

(i) L expresses a λ-term.

(ii) $\vdash_{\mathbf{Reg}^+_{\mathsf{letrec}}} ()\ L$.

Proof. For the proof of both directions of the equivalence, let $L \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$.

For showing the implication (i) $\Rightarrow$ (ii), we argue indirectly, and therefore assume that $()\ L$ is not derivable in $\mathbf{Reg}^+_{\mathsf{letrec}}$. Then, while $()\ L$ is not derivable in $\mathbf{Reg}^+_{\mathsf{letrec}}$, there is, by proposition 1.6.25 (i), a derivation $\mathcal{D}$ in $\mathbf{ug} - \mathbf{Reg}^+_{\mathsf{letrec}}$ that is not guarded. It follows by lemma 1.6.24, and in particular due to lemma 1.6.24 (iv), that there is a scope$^+$-delimiting strategy $\mathbb{S}_{\mathcal{D}}$ for Reg^+ such that L is not $\mathbb{S}_{\mathcal{D}}$-productive. Then it follows by lemma 1.5.20, using (i) $\Rightarrow$ (iv) there, that L does not unfold to a λ-term.

For showing the implication (ii) $\Rightarrow$ (i), let $\mathcal{D}$ be a closed derivation in $\mathbf{Reg}^+_{\mathsf{letrec}}$ with conclusion $()\ L$. It follows by lemma 1.6.24 that there is a scope$^+$-delimiting strategy $\mathbb{S}$ for Reg such that L is $\mathbb{S}$-productive. Then lemma 1.5.20 implies that L unfolds to a λ-term. □

§ 1.6.27 (soundness and completeness of $\mathbf{Reg}_{\mathsf{letrec}}$)**.** Also the proof system $\mathbf{Reg}_{\mathsf{letrec}}$ can be shown to be sound and complete for the property of $\lambda_{\mathsf{letrec}}$-terms to unfold to λ-terms. To establish this in analogy with the route of proof we pursued here, a CRS $\boldsymbol{Parse}$ similar to $\boldsymbol{Parse}^{+}$ (see definition 1.4.52) could be defined by replacing the rule $\varrho^{\mathsf{S}}_{\mathrm{parse}^{+}}$ by a rule $\varrho^{\mathrm{del}}_{\mathrm{parse}^{+}}$ that can compress more abstraction prefixes, similar as the rule $\varrho^{\mathrm{del}}_{\mathrm{reg}}$ of $\boldsymbol{Reg}$ can compress more abstraction prefixes than the rule $\varrho^{\mathsf{S}}_{\mathrm{reg}^{+}}$ of $\boldsymbol{Reg}^{+}$. Then furthermore also a lemma analogous to lemma 1.5.20 can be formulated, proved, and used in a similar way.

We now arrive at a theorem that states one direction of our main characterisation result (theorem 1.8.13 in section 1.8) that will link $\lambda_{\mathsf{letrec}}$-expressibility to strong regularity of λ-terms.

Theorem 1.6.28. Every $\lambda_{\mathsf{letrec}}$-expressible λ-term is strongly regular.

Proof. Let M be a λ-term that is expressible by a $\lambda_{\mathsf{letrec}}$-term L, that is, $L \twoheadrightarrow^{\omega}_{\triangledown} M$ holds. Then by theorem 1.6.26 there exists a closed derivation $\mathcal{D}$ in $\mathbf{Reg}^{+}_{\mathsf{letrec}}$ with conclusion $() \; L$. Now lemma 1.6.24 guarantees a scope^{+}-delimiting strategy $\mathbb{S}_{\mathcal{D}}$ for $Reg^{+}_{\mathsf{letrec}}$ such that L is $\mathbb{S}_{\mathcal{D}}$-regular. Then lemma 1.5.17 gives a scope^{+}-delimiting strategy $\check{\mathbb{S}}_{\mathcal{D}}$ for $Reg^{+}_{\mathsf{letrec}}$ such that $M = [\![L]\!]_{\boldsymbol{\lambda}}$ is $\check{\mathbb{S}}_{\mathcal{D}}$-regular. It follows that M is strongly regular. □

1.7 Binding–Capturing Chains

§ 1.7.1 (overview)**.** In this section we develop a characterisation for strong regularity of λ-terms by means of the 'binding–capturing chains' occurring in a term. This concept is related to the notions of scope and scope^{+} as explained informally in § 1.4.1. Binding–capturing chains occur whereever scopes overlap; and they are fully contained within scopes^{+}. First we give definitions for the concepts involved: binding, capturing, and binding–capturing chains. Then we show that strong regularity of regular λ-terms is equivalent to the absence of infinite binding–capturing chains.

§ 1.7.2 (binding and capturing)**.** We will define binding and capturing as relations on the positions of a λ-term. Binding relates an abstraction with the occurrences of the variable it binds. If p is the position of an abstraction $(\lambda x. \ldots)$ that abstracts over x and q is the position of an occurrence of x that is bound by the abstraction, then we will write $p \multimap q$ and say that p 'binds' q.

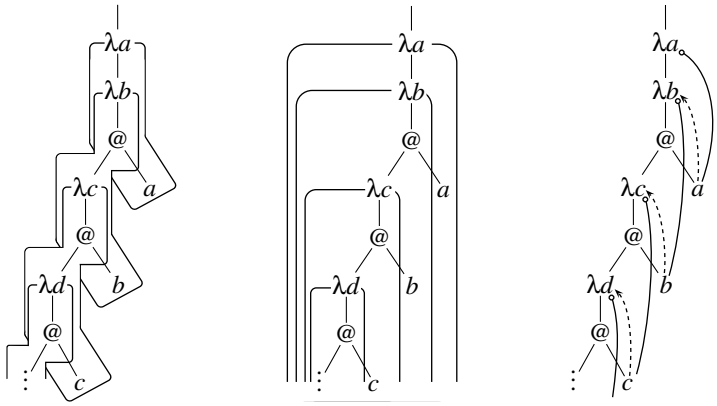

Figure 1.22. The term graph of the term in example 1.2.2 with its overlapping scopes (left), its nested scopes$^+$ (middle), and with its binding ($\circ$—) and capturing (-$\rightarrow$) links (right).

Capturing relates an abstraction with variables that occur freely underneath the abstraction. If p is the position of an abstraction, and $q > p$ is the position of a variable that is free in the entire subterm at position p, then we will write $p \leftarrow q$ and say that p 'captures' q. See fig. 1.22 for an illustration of these concepts.

§ 1.7.3 (positions in iCRS terms)**.** When we speak of positions in λ-terms (and thus iCRS terms) we act on the assumption that positions on iCRS terms are an established concept as for example in [31]. Note, however, that we deviate slightly from the scheme there in addressing the arguments of an app by 0 and 1 instead of 1 and 2.

§ 1.7.4 (binding–capturing chains in the literature)**.** Binding–capturing chains have been used in [15] to study α-avoiding rewrite sequences in a rewrite calculus for μ-unfolding. They originate from the notion of 'gripping' due to [40], and from techniques developed in [46] concerning the notion of 'holding' of redexes (which is shown there as being 'parting' for CRSs, that is, never relating two residuals of the same redex).

We now define 'binding' and 'capturing' formally as binary relations on the set of positions of λ-terms.

Definition 1.7.5 (binding, capturing)**.** For every $M \in \mathrm{Ter}^{\infty}(\boldsymbol{\lambda})$ we define the binary relations $\multimap$ and $\twoheadrightarrow$ on the set $\mathrm{Pos}(M) \subseteq \mathbb{N}^*$ of positions of M:

(i) We say that a binder at position p *binds* a variable occurrence at position q, symbolically $p \multimap q$, if p is a binder position, and q a variable position in M, and the binder at position p binds the variable occurrence at position q.

(ii) We say that a variable occurrence at position q *is captured by* a binder at position p, symbolically $q \twoheadrightarrow p$ (and that a binder at position p *captures* a variable occurrence at position q, symbolically $p \twoheadleftarrow q$), if q is a variable position and $p < q$ a binder position in M, and there is no binder position q_0 in M with $p \leq q_0$ and $q_0 \multimap q$.

Definition 1.7.6 (binding–capturing chain)**.** Let M be a λ-term. A finite or infinite sequence $\langle p_0, p_1, p_2, \ldots \rangle$ in $\mathrm{Pos}(M)$ is called a *binding–capturing chain in* M if p_0 is the position of an abstraction in M, and the positions in the sequence are alternatingly linked via binding and capturing, starting with a binding: $p_0 \multimap p_1 \twoheadrightarrow p_2 \multimap \ldots$.

§ 1.7.7. Binding–capturing chains are closely related to the notion of scope and scope$^+$. In order to establish this, we first give precise definitions of the notions of scope and scope$^+$ in terms of an 'in-scope' rewrite relation on the positions of a λ-term: While the scope of a binder position p is the set of positions between p and variable positions bound at p (the positions directly reachable by a single 'in-scope' step), the scope$^+$ of p is the set of positions reachable by a finite number of successive 'in-scope' steps.

Notation 1.7.8 (binder positions)**.** In the following paragraphs for a given position p in some λ-term, we write $\mathrm{bp}(p)$ for the proposition 'p is a binder position'.

Definition 1.7.9 (scope and scope$^+$)**.** Let M be a λ-term. On the set $\mathrm{Pos}(M)$ of M, the *in-scope* relation $\to_{\mathrm{sc}}$ (for M) is defined by:

$$p \to_{\mathrm{sc}} q \iff \mathrm{bp}(p) \ \wedge \ \exists p' \in \mathrm{Pos}(M) \ \ p \multimap^{=} p' \ \wedge \ p \leq q \leq p'$$

where $\multimap^{=}$ denotes the reflexive closure of the binding relation $\multimap$.

For every binder position $p \in \mathrm{Pos}(M)$, the *scope of* p *in* M and the *scope*$^+$ *of* p *in* M are defined as the following sets of positions in M:

$$\mathrm{scope}_M(p) := \{q \in \mathrm{Pos}(M) \mid p \to_{\mathrm{sc}} q\}$$
$$\mathrm{scope}^+_M(p) := \{q \in \mathrm{Pos}(M) \mid p \to^+_{\mathrm{sc}} q\}$$

(Note that the scopes and scopes$^+$ of non-binder positions are empty sets.)

The following proposition establishes that binding–capturing chains starting at a binder position p span the positions of the scope$^+$ of p.

Proposition 1.7.10 (binding, capturing, and scope/scope$^+$)**.** Let M be a λ-term and $p \in \mathrm{Pos}(M)$ be a binder position. Then for all positions $q \in \mathrm{Pos}(M)$ the following statements hold:

(i) $p \to_{\mathrm{sc}} q \;\wedge\; \mathrm{bp}(q) \iff p = q \;\vee\; p \circ\!\!- \cdot \to q$

(ii) $p \to_{\mathrm{sc}}^{+} q \iff \exists p' \in \mathrm{Pos}(M)\;\; p\ (\circ\!\!- \cdot \to)^* \cdot \circ\!\!-^{=} p' \wedge p \le q \le p'$

(iii) $\mathrm{scope}_M(p) = \left\{ q \in \mathrm{Pos}(M) \;\middle|\; \exists p' \in \mathrm{Pos}(M)\;\; \begin{array}{l} p \circ\!\!-^{=} p' \\ \wedge\, p \le q \le p' \end{array} \right\}$

(iv) $\mathrm{scope}_M^{+}(p) = \left\{ q \in \mathrm{Pos}(M) \;\middle|\; \exists p' \in \mathrm{Pos}(M)\;\; \begin{array}{l} p\ (\circ\!\!- \cdot \to)^* \cdot \circ\!\!-^{=} p' \\ \wedge\, p \le q \le p' \end{array} \right\}$

Conversely, every position that is covered by a binding–capturing chain starting at a binding position p is in the scope$^+$ of p:

Proposition 1.7.11. Let $\langle p_0, p_1, p_2, \ldots \rangle$ be a binding–capturing chain in a λ-term M. Then it holds that $p_0 < p_2 < \ldots$, and $p_0 < p_1$, $p_2 < p_3$, ... Furthermore, for all q with $p_0 \le q \le p_n$ for some $n \in \mathbb{N}$ with p_n being a position on the chain it holds that $q \in \mathrm{scope}_M^{+}(p_0)$.

§ 1.7.12 (position-annotated variants $\boldsymbol{Reg_{pos}}$ and $\boldsymbol{Reg^{+}_{pos}}$)**.** In order to study the relationship between rewrite sequences in $\boldsymbol{Reg^{+}}$ and binding–capturing chains we first introduce a position-annotated variant of $\boldsymbol{Reg^{+}}$. The idea is, that when a prefixed term $(\lambda y_1 \ldots y_n)\ N$ is obtained as a generated subterm of a λ-term M by a $\to_{\mathrm{reg}}$ or $\to_{\mathrm{reg}^+}$ rewrite sequence τ on $()\ M$, then in the position-annotated rewrite system a prefixed term $(\lambda y_1, \ldots, y_n)^{q}_{p_1 \ldots p_n}\ N$ is obtained by an annotated version $\hat{\tau}$ of the rewrite sequence τ such that: the positions $p_1, \ldots, p_n$ are the positions in (the original λ-term) M from which the bindings λy_1, ..., λy_n in the abstraction prefix stem from; and q is the position in M of the body N of the subterm generated by τ.

§ 1.7.13 (position-annotated decomposition in informal notation)**.** On $\mathrm{Ter}^{\infty}(\boldsymbol{(\lambda)})$ we consider the following rewrite rules:

$$\varrho^{@_i}_{pos}:\ (\lambda x_1 \dots x_n)^{q}_{p_1,\dots,p_n}\ M_0\ M_1 \to (\lambda x_1 \dots x_n)^{qi}_{p_1,\dots,p_n}\ M_i \quad (i \in \{0,1\})$$

$$\varrho^{\lambda}_{pos}:\ (\lambda x_1 \dots x_n)^{q}_{p_1,\dots,p_n}\ \lambda y.\, M_0 \to (\lambda x_1 \dots x_n y)^{q00}_{p_1,\dots,p_n,q}\ M_0$$

$$\varrho^{\mathsf{S}}_{pos}:\ (\lambda x_1 \dots x_{n+1})^{q}_{p_1,\dots,p_{n+1}}\ M_0 \to (\lambda x_1 \dots x_n)^{q}_{p_1,\dots,p_n}\ M_0$$
(if the binding λx_{n+1} is vacuous)

$$\varrho^{\mathrm{del}}_{pos}:\ (\lambda x_1 \dots x_{n+1})^{q}_{p_1,\dots,p_{n+1}}\ M_0 \to (\lambda x_1 \dots x_{i-1} x_{i+1} \dots x_{n+1})^{q}_{p_1,\dots p_{i-1},p_{i+1}\dots p_{n+1}}\ M_0$$
(if the binding λx_i is vacuous)

Note that in the rule ϱ^{λ}_{pos} the position changes from q to $q00$, because of the underlying CRS notation for terms in $\boldsymbol{(\lambda)}$: when a term $\mathsf{abs}([y]\,M_0)$ starts at position q, then the CRS-binding is at position $q0$, and the body M_0 starts at position $q00$.

Definition 1.7.14 (position-annotated decomposition rewrite systems $\boldsymbol{Reg^{-}_{pos}}$, $\boldsymbol{Reg_{pos}}$, $\boldsymbol{Reg^{+}_{pos}}$)**.** The CRS signature for $\boldsymbol{(\lambda)_{pos}}$, the λ-calculus with position-annotated abstraction prefixes is given by $\Sigma^{(\lambda)}_{\mathrm{pos}} = \Sigma^{\lambda} \cup \left\{ \mathsf{pre}^{q}_{\langle p_1,\dots,p_n\rangle} \;\middle|\; p_1,\dots,p_n,q \in \{0,1\}^* \right\}$ where all of the function symbols $\mathsf{pre}^{q}_{\langle p_1,\dots,p_n\rangle}$ are unary. We consider the following CRS rule schemes over $\Sigma^{(\lambda)}_{\mathrm{pos}}$:

$$\varrho^{@_i}_{pos}:\ \mathsf{pre}^{q}_{\langle p_1,\dots,p_n\rangle}([x_1 \dots x_n]\,\mathsf{app}(M_0,\, M_1)) \to \mathsf{pre}^{qi}_{\langle p_1,\dots,p_n\rangle}([x_1 \dots x_n]\,M_i) \quad (i \in \{0,1\})$$

$$\varrho^{\lambda}_{pos}:\ \mathsf{pre}^{q}_{\langle p_1,\dots,p_n\rangle}([x_1 \dots x_n]\,\mathsf{abs}([y]\,M_0)) \to \mathsf{pre}^{q00}_{\langle p_1,\dots,p_n,q\rangle}([x_1 \dots x_n y]\,M_0)$$

$$\varrho^{\mathsf{S}}_{pos}:\ \mathsf{pre}^{q}_{\langle p_1,\dots,p_{n+1}\rangle}([x_1 \dots x_{n+1}]\,M_0) \to \mathsf{pre}^{q}_{\langle p_1,\dots,p_n\rangle}([x_1 \dots x_n]\,M_0)$$

$$\varrho^{\mathrm{del}}_{pos}:\ \mathsf{pre}^{q}_{\langle p_1,\dots,p_{n+1}\rangle}([x_1 \dots x_{n+1}]\,M_0) \to \mathsf{pre}^{q}_{\langle p_1,\dots p_{i-1},p_{i+1}\dots p_{n+1}\rangle}([x_1 \dots x_{i-1} x_{i+1} \dots x_{n+1}]\,M_0)$$

By $\boldsymbol{Reg^{-}_{pos}}$ we denote the CRS with the rules $\varrho^{@_i}_{pos}$ and ϱ^{λ}_{pos}. By $\boldsymbol{Reg_{pos}}/\boldsymbol{Reg^{+}_{pos}}$ we denote the CRS consisting of the rules $\varrho^{@_i}_{pos}$, ϱ^{λ}_{pos}, and $\varrho^{\mathrm{del}}_{pos}/\varrho^{\mathsf{S}}_{pos}$.

By Reg^{-}_{pos}, Reg_{pos} and Reg^{+}_{pos} we denote the ARSs induced by the iCRSs derived from $\boldsymbol{Reg^{-}_{pos}}$, $\boldsymbol{Reg_{pos}}$, $\boldsymbol{Reg^{+}}$, restricted to position-annotated terms in $\mathrm{Ter}^{\infty}(\boldsymbol{(\lambda)})$.

Proposition 1.7.15 (position-annotated rewrite sequences)**.** It holds:

(i) Every rewrite sequence

$$\tau:\ (\lambda\vec{x}_0)\ M_0 \to (\lambda\vec{x}_1)\ M_1 \to \ldots \to (\lambda\vec{x}_n)\ M_n \tag{1.2}$$

in Reg^-, Reg, or Reg^+ can be transformed (lifted) step by step, for a given $q_0 \in \mathbb{N}^*$ and $\vec{p}_0 \in \vec{\mathbb{N}}^*$ with $|\vec{p}_0| = |\vec{x}_0|$, by adding these and appropriate further position annotations $q_1, \ldots, q_n \in \mathbb{N}^*$ and $\vec{p}_1, \ldots, \vec{p}_n \in \vec{\mathbb{N}}^*$, to a rewrite sequence:

$$\hat{\tau}:\ (\lambda\vec{x}_0)^{q_0}_{\vec{p}_0}\ M_0 \to (\lambda\vec{x}_1)^{q_1}_{\vec{p}_1}\ M_1 \to \ldots \to (\lambda\vec{x}_n)^{q_n}_{\vec{p}_n}\ M_n \tag{1.3}$$

in Reg^-_{pos}, Reg_{pos}, or Reg^+_{pos}, accordingly, such that the result of dropping the position annotations in the prefix of $\hat{\tau}$ is again τ.

(ii) Conversely, every rewrite sequence ξ in Reg^-_{pos}, Reg_{pos}, or Reg^+_{pos} of the form (1.3) can be transformed step by step, by dropping the position annotations in the prefix, to a rewrite sequence $\check{\xi}$ of the form (1.3) in Reg^-, Reg, or Reg^+, respectively.

The transformations in (i) and (ii) preserve eagerness/laziness of rewrite sequences.

The proposition below characterises the binding relation $\circ\!\!-$ and the capturing relation $-\!\!\circ$ on the positions of an infinite term M with the help of rewrite sequences with respect to $\to_{\mathrm{reg}^-}$ on $()^\epsilon\ M$ in $\boldsymbol{Reg^-_{pos}}$ down to 'variable occurrences' $(\lambda\vec{x})^q_{\vec{p}}\ x_i$ in M.

Proposition 1.7.16 (binding, capturing, and $\boldsymbol{Reg^-_{pos}}$ rewrite sequences)**.** For all $M \in \mathrm{Ter}^\infty(\boldsymbol{\lambda})$ and positions $p, q \in \mathrm{Pos}(M)$ it holds:

$$p \circ\!\!- q \Leftrightarrow \begin{array}{l}\text{there is a rewrite sequence } ()^\epsilon_{\langle\rangle}\ M \twoheadrightarrow_{\mathrm{reg}^-} (\lambda x_1 \ldots x_n)^q_{p_1,\ldots,p_n}\ x_i \\ \text{with } x_1 \ldots x_n \text{ distinct, } i \in \{1,\ldots,n\}, \text{ such that } p = p_i\end{array}$$

$$q -\!\!\circ p \Leftrightarrow \begin{array}{l}\text{there is a rewrite sequence } ()^\epsilon_{\langle\rangle}\ M \twoheadrightarrow_{\mathrm{reg}^-} (\lambda x_1 \ldots x_n)^q_{p_1,\ldots,p_n}\ x_i \\ \text{with } x_1 \ldots x_n \text{ distinct, } i \in \{1,\ldots,n\}, \text{ such that } p \in \{p_{i+1},\ldots,p_n\}\end{array}$$

The following lemmas describe the close relationship between, on the one hand, binding–capturing chains in a λ-term M, and on the other hand, $\to_{\mathrm{reg}^+}$-rewrite-sequences on $()^\epsilon\ M$ in $\boldsymbol{Reg^+_{pos}}$ that are guided by the eager scope$^+$-delimiting strategy.

Lemma 1.7.17 (binding–capturing chains)**.** For all $M \in \mathrm{Ter}^{\infty}(\boldsymbol{\lambda})$ the following statements hold:

(i) If there is a rewriting sequence of the form

$$\begin{aligned}()^{\epsilon}\ M &\twoheadrightarrow_{\mathbb{S}^{+}_{\mathrm{eag}}} (\lambda x_0 \ldots x_{n_1})^{q}_{p_0,\ldots,p_{n_1}}\ N \\ &\twoheadrightarrow_{\mathbb{S}^{+}_{\mathrm{eag}}} (\lambda x_0 \ldots x_{n_1} \ldots x_{n_2})^{q'}_{p_0,\ldots,p_{n_1},\ldots,p_{n_2}}\ O\end{aligned}$$

then there exist $q_{n_1+1}, \ldots, q_{n_2} \in \mathrm{Pos}(M)$ such that $p_{n_1} \multimap q_{n_1+1} \dashrightarrow p_{n_1+1} \multimap \ldots \multimap q_{n_2} \dashrightarrow p_{n_2}$.

(ii) If $p_0 \multimap q_1 \dashrightarrow p_1 \multimap \ldots \multimap q_n \dashrightarrow p_n$ is a binding–capturing chain in M, then there exist $r_0, \ldots, r_m, s \in \mathrm{Pos}(M)$ with $m \geq n$ such that $()^{\epsilon}\ M \twoheadrightarrow_{\mathbb{S}^{+}_{\mathrm{eag}}} (\lambda x_0 \ldots x_m)^{s}_{r_0,\ldots,r_m}\ N$ and furthermore $p_0, \ldots, p_n \in \{r_0, \ldots, r_m\}$ such that $p_0 < p_1 < \ldots < p_n = r_m$.

Lemma 1.7.18 (length of binding–capturing chains)**.** Let M be a λ-term such that $()\ M \twoheadrightarrow_{\mathbb{S}^{+}_{\mathrm{eag}}} (\lambda x_0 \ldots x_n)\ N$. Then M contains a binding–capturing chain of length n.

Proof. By proposition 1.7.15 (i), the assumed rewrite sequence $()\ M \twoheadrightarrow_{\mathbb{S}^{+}_{\mathrm{eag}}} (\lambda x_0 \ldots x_n)\ N$ in Reg^{+} can be lifted to a rewrite sequence $()^{\epsilon}\ M \to_{\mathbb{S}^{+}_{\mathrm{eag}}} (\lambda x_0 \ldots x_n)^{q}_{p_0,\ldots,p_n}\ N$ in Reg^{+}_{pos}. Then by lemma 1.7.17 (i), there exists a binding–capturing chain of length n. □

§ 1.7.19 (binding–capturing chains and scope/scope^{+})**.** The notion of scope and scope^{+} helps to understand the relationship between binding–capturing chains and rewrite sequences in $\boldsymbol{Reg^{+}}$. A binding–capturing chain corresponds to the overlap of scopes, or in other words the nesting of scopes^{+}. An infinite binding–capturing chain thus corresponds to a infinitely deep nesting of scopes^{+} and therefore to an unrestricted growth of the prefix in certain rewriting sequences in $\boldsymbol{Reg^{+}}$.

Lemma 1.7.20 (infinite binding–capturing chains)**.** Let M be a λ-term, and let τ be an infinite rewrite sequence in Reg w.r.t. the eager scope^{+}-delimiting strategy $\mathbb{S}^{+}_{\mathrm{eag}}$:

$$\tau\ :\ ()\ M = (\lambda \vec{x}_0)\ M_0 \to_{\mathbb{S}^{+}_{\mathrm{eag}}} (\lambda \vec{x}_1)\ M_1 \to_{\mathbb{S}^{+}_{\mathrm{eag}}} \ldots \to_{\mathbb{S}^{+}_{\mathrm{eag}}} (\lambda \vec{x}_i)\ M_i \to_{\mathbb{S}^{+}_{\mathrm{eag}}} \ldots$$

Furthermore suppose that for $\mathrm{p} : \mathbb{N} \to \mathbb{N}$, $i \mapsto \mathrm{p}(i) := |\vec{x}_i|$, the prefix length function associated with τ, there exists a lower bound $\mathrm{lb} : \mathbb{N} \to \mathbb{N}$ such that

lb is non-decreasing, and $\lim_{n\to\infty} \mathrm{lb}(n) = \infty$. Then there exists an infinite binding–capturing chain in M.

Proof. Let M, τ, p, lb as above. By proposition 1.7.15 (i), τ can be lifted to position annotated counterpart

$$\hat{\tau} : ()^{\epsilon}\, M = (\lambda\vec{x}_0)^{\epsilon}\, M_0 \to_{\mathbb{S}^+_{\mathrm{eag}}} (\lambda\vec{x}_1)^{q_1}_{\vec{p}_1}\, M_1 \to_{\mathbb{S}^+_{\mathrm{eag}}} \ldots \to_{\mathbb{S}^+_{\mathrm{eag}}} (\lambda\vec{x}_i)^{q_i}_{\vec{p}_i}\, M_i \to_{\mathbb{S}^+_{\mathrm{eag}}} \ldots$$

where, for all $i \in \mathbb{N}$, q_i are positions and $\vec{p}_i = \langle p_1, \ldots, p_{m_i}\rangle$ vectors of positions, with $m_i \in \mathbb{N}$. Next we define the function:

$$\mathrm{st} : \mathbb{N} \to \mathbb{N}, \quad l \mapsto \mathrm{st}(l) := \min\{i \mid \mathrm{lb}(i) \geq l\}$$

which is well-defined, since $\lim_{n\to\infty} \mathrm{lb}(n) = \infty$. It describes a prefix stabilisation property: for every $l \in \mathbb{N}$, it gives the first index $i = \mathrm{st}(l)$ with the property that the prefix of $(\lambda\vec{x}_i)\, M_i$ contains more than l abstractions, and (since lb is non-decreasing, and a lower bound for p) that from i onward the l-th abstraction never disappears again, for $j \geq i$, in terms $(\lambda\vec{x}_j)\, M_j$ that follow in τ as well as in $\hat{\tau}$. Furthermore, st is non-decreasing, as an easy consequence of its definition, and unbounded: if st were bounded by $M \in \mathbb{N}$, then $\forall l \in \mathbb{N}\ \exists i \in \mathbb{N}\ i \leq M \wedge \mathrm{lb}(i) \geq l$ would follow, which cannot be the case since $\{\mathrm{lb}(0), \ldots, \mathrm{lb}(M)\}$ is a finite set. By non-decreasingness and unboundedness it also follows that $\lim_{n\to\infty} \mathrm{st}(n) = \infty$.

So when the rewrite sequence $\hat{\tau}$ is split into segments indicated in

$$\begin{aligned}()^{\epsilon}\, M \twoheadrightarrow_{\mathbb{S}^+_{\mathrm{eag}}} \ldots \twoheadrightarrow_{\mathbb{S}^+_{\mathrm{eag}}} & \left(\lambda\vec{x}_{\mathrm{st}(i)}\right)^{q_{\mathrm{st}(i)}}_{\vec{p}_{\mathrm{st}(i)}}\, M_{\mathrm{st}(i)} \\ & \twoheadrightarrow_{\mathbb{S}^+_{\mathrm{eag}}} \left(\lambda\vec{x}_{\mathrm{st}(i+1)}\right)^{q_{\mathrm{st}(i+1)}}_{\vec{p}_{\mathrm{st}(i+1)}}\, M_{\mathrm{st}(i+1)} \twoheadrightarrow_{\mathbb{S}^+_{\mathrm{eag}}} \ldots\end{aligned}$$

then it follows that all terms of the sequence after $\left(\lambda\vec{x}_{\mathrm{st}(i)}\right)^{q_{\mathrm{st}(i)}}_{\vec{p}_{\mathrm{st}(i)}}\, M_{\mathrm{st}(i)}$ have an abstraction prefix of length greater or equal to i, for all $i \in \mathbb{N}$.

Now note that in a step

$$(\lambda\vec{x})^{q}_{\langle p_1,\ldots,p_n\rangle}\, N \to (\lambda\vec{x}')^{q'}_{\left\langle p'_1,\ldots,p'_{n'}\right\rangle}\, N'$$

in Reg_{pos} that does not shorten the abstraction prefix it holds that $n \leq n'$ and $p'_1 = p_1$, $\ldots$, $p'_n = p_n$, that is, positions in the vector in the subscript of the abstraction prefix are preserved. As a consequence it follows for the rewrite sequence $\hat{\tau}$ that for all $i \in \mathbb{N}$ and $j > i$ the position vector $\vec{p}_{\mathrm{st}(j)}$ in the term $\left(\lambda\vec{x}_{\mathrm{st}(j)}\right)^{q_{\mathrm{st}(j)}}_{\vec{p}_{\mathrm{st}(j)}}\, M_{\mathrm{st}(j)}$ is of the form

$$\vec{p}_{\mathrm{st}(j)} = \left\langle p_{1,\mathrm{st}(i)}, \ldots, p_{i,\mathrm{st}(i)}, p_{j,\mathrm{st}(j)}, \ldots, p_{j,m_j}\right\rangle.$$

This implies furthermore that for all $i \in \mathbb{N}$:

$$\vec{p}_{\mathrm{st}(i)} = \left\langle p_{1,\mathrm{st}(1)}, p_{2,\mathrm{st}(2)}, \ldots, p_{i,\mathrm{st}(i)}, \ldots, p_{i,m_i} \right\rangle.$$

Then lemma 1.7.17 (i), implies the existence of positions $q_2, q_3, \ldots$ such that

$$p_{1,\mathrm{st}(1)} \multimap q_2 \dashrightarrow p_{2,\mathrm{st}(2)} \multimap q_3 \dashrightarrow \ldots \dashrightarrow p_{i,\mathrm{st}(i)} \multimap q_{i+1} \dashrightarrow p_{i+1,\mathrm{st}(i+1)} \multimap \cdots$$

and thereby, an infinite binding–capturing chain in M. □

Now we formulate and prove the main theorem of this section, which applies the concept of binding–capturing chain to pin down, among all λ-terms that are regular, those that are strongly regular.

Theorem 1.7.21 (binding–capturing chains and strong regularity)**.** A regular λ-term is strongly regular if and only if it contains only finite binding–capturing chains.

And taking into account proposition 1.4.43 (i), we obtain:

Corollary 1.7.22. A λ-term is strongly regular if and only if it is regular and contains only finite binding–capturing chains.

Proof of theorem 1.7.21. Let M be a λ-term that is regular.

For showing the implication "$\Rightarrow$", we assume that M is also strongly regular. Then there exists a scope$^+$-delimiting strategy $\mathbb{S}$ such that $\mathrm{ST}_{\mathbb{S}}(M)$ is finite. By proposition 1.4.44 (i) it follows that then also $\mathrm{ST}_{\mathbb{S}^+_{\mathrm{eag}}}(M)$ is finite for the eager scope$^+$-delimiting strategy $\mathbb{S}^+_{\mathrm{eag}}$ in Reg^+. Now let n be the longest abstraction prefix of a term in $\mathrm{ST}_{\mathbb{S}^+_{\mathrm{eag}}}(M)$. Then it follows by lemma 1.7.18 that the length of every binding–capturing chain in M is bounded by $n-1$. Hence M only contains finite binding–capturing chains.

In the rest of this proof, we establish the implication "$\Leftarrow$" in the statement of the theorem. For this we argue indirectly: assuming that M is not strongly regular, we show the existence of an infinite binding–capturing chain in M.

So suppose that M is not strongly regular. Then for all scope$^+$-delimiting strategies $\mathbb{S}$ in Reg^+ it holds that $\mathrm{ST}_{\mathbb{S}}(M)$ is infinite. This means that in particular $\mathrm{ST}_{\mathbb{S}^+_{\mathrm{eag}}}(M)$ is infinite for the eager scope-delimiting strategy $\mathbb{S}^+_{\mathrm{eag}}$ on Reg^+. It follows that the number of $\twoheadrightarrow_{\mathbb{S}^+_{\mathrm{eag}}}$-reducts, and hence the generated sub-ARS $\left(() \; M \twoheadrightarrow_{\mathbb{S}^+_{\mathrm{eag}}}\right)$ of $() \; M$ in Reg^+ is infinite. Since $\twoheadrightarrow_{\mathbb{S}^+_{\mathrm{eag}}}$ on Reg^+ has

branching degree ≤ 2 (branching actually only happens at sources of $\to_{@_i}$-steps), it follows by Kőnig's Lemma that there exists an infinite rewrite sequence:

$$\tau : \ () \ M = (\lambda \vec{x}_0) \ M_0 \to_{\mathbb{S}^+_{\mathrm{eag}}} \ldots \to_{\mathbb{S}^+_{\mathrm{eag}}} (\lambda \vec{x}_i) \ M_i \to_{\mathbb{S}^+_{\mathrm{eag}}} \ldots \tag{1.4}$$

in Reg^+ that passes through distinct terms. By lemma 1.4.29 (i), this rewrite sequence projects to a rewrite sequence:

$$\check{\tau} : \ () \ M = (\lambda \vec{x}'_0) \ M_0 \twoheadrightarrow_{\mathbb{S}_{\mathrm{eag}}} \ldots \twoheadrightarrow_{\mathbb{S}_{\mathrm{eag}}} (\lambda \vec{x}'_i) \ M_i \twoheadrightarrow_{\mathbb{S}_{\mathrm{eag}}} \ldots \tag{1.5}$$

in *Reg* in the sense that:

$$(\lambda \vec{x}_i) \ M_i \twoheadrightarrow_{\mathrm{del}} (\lambda \vec{x}'_i) \ M_i \quad (\text{for all } i \in \mathbb{N}). \tag{1.6}$$

Note that, in the terms, the projection merely shortens the length of the abstraction prefix. Since M is regular, $\mathrm{ST}_{\mathbb{S}_{\mathrm{eag}}}(M)$ is finite by proposition 1.4.44 (i), and hence it follows that only finitely many terms occur in $\check{\tau}$.

Now we will use this contrast with τ, and the fact that the terms of τ project to terms in $\check{\tau}$ via $\twoheadrightarrow_{\mathrm{del}}$-prefix compression rewrite sequences, to show that the prefix lengths in terms of τ are unbounded, and stronger still, that these lengths actually tend to infinity. More precisely, we show the following:

$$\forall l \in \mathbb{N} \ \exists i_0 \in \mathbb{N} \ \forall i \geq i_0 \ \ |\vec{x}_i| \geq l \tag{1.7}$$

Suppose that this statement does not hold. Then there exists $l_0 \in \mathbb{N}$ such that $|\vec{x}_i| < l_0$ for infinitely many $i \in \mathbb{N}$. This means that there is an increasing sequence $i_0 < i_1 < i_2 < i_3 < \ldots$ in $\mathbb{N}$ such that:

$$S \coloneqq \left\{ \left(\lambda \vec{x}_{i_j}\right) M_{i_j} \ \middle| \ j \in \mathbb{N} \right\} \text{ is infinite} \tag{1.8}$$

$$\text{and for all } \left(\lambda \vec{x}_{i_j}\right) M_{i_j} \in S \ \ \left|\vec{x}_{i_j}\right| < l_0 \tag{1.9}$$

(S is infinite since the terms in τ are distinct). On the other hand we have:

$$T \coloneqq \left\{ \left(\lambda \vec{x}'_{i_j}\right) M_{i_j} \ \middle| \ j \in \mathbb{N} \right\} \ \subseteq \ \mathrm{ST}_{\mathbb{S}_{\mathrm{eag}}}(M) \text{ is finite} \tag{1.10}$$

because M is regular. However, since every term in S has a $\twoheadrightarrow_{\mathrm{del}}$-reduct in T due to (1.6), as well as an abstraction prefix of a length bounded by l_0, it follows by proposition 1.4.26 (ii), that S also has to be finite, conflicting with (1.8). We have reached a contradiction, and thereby established (1.7).

Now we are able to define a lower bound on the lengths of the prefixes in τ that fulfils the requirements of lemma 1.7.20. We define the function:

$$\mathrm{lb} : \mathbb{N} \to \mathbb{N}, \quad n \mapsto \mathrm{lb}(n) \coloneqq \min\{|\vec{x}_{n'}| \mid n' \geq n\}$$

Its definition guarantees that lb is a lower bound on the prefix lengths in τ, and that lb is non-decreasing. Furthermore also $\lim_{n\to\infty} \mathrm{lb}(n) = \infty$ follows by non-decreasingness, in addition to unboundedness of lb: for arbitrary $l \in \mathbb{N}$, by (1.7) there exists $n_0 \in \mathbb{N}$ such that $|\vec{x}_n| \geq l$ holds for all $n \in \mathbb{N}$, $n \geq n_0$; this entails $\mathrm{lb}(n_0) \geq l$.

Now since M, τ, together with lb as defined above, satisfy the assumptions of lemma 1.7.20, this lemma can be applied, yielding an infinite binding–capturing chain in M. □

Example 1.7.23 (infinite binding–capturing chain)**.** The infinite λ-term from example 1.2.2 with a representation as a higher-order recursive program scheme in example 1.4.46, which was recognised there to be regular but not strongly regular, possesses an infinite binding–capturing chain as indicated on the right in fig. 1.22.

1.8 Expressibility by terms in $\boldsymbol{\lambda}_{\mathsf{letrec}}$

§ 1.8.1 (overview)**.** In this section we finish the proof of our main characterisation result: we prove that every strongly regular λ-term is $\lambda_{\mathsf{letrec}}$-expressible. For this purpose we introduce an annotated variant of one of the proof systems for strongly regular λ-terms. We show that every closed derivation in $\mathbf{Reg}_{\mathbf{0}}^{+}$ with conclusion $()\ M$, which witnesses that M is strongly regular, can be annotated, by adding appropriate $\lambda_{\mathsf{letrec}}$-terms to each prefixed term in the derivation, into a derivation in the annotated system with conclusion $()\ L : M$ such that the $\lambda_{\mathsf{letrec}}$-term annotation L expresses the λ-term M. We show the correctness of this construction by transforming the derivation in the annotated proof system into a derivation in the proof system $\mathbf{Reg}_{=}^{+}$ with conclusion $()\ [\![L]\!]_{\boldsymbol{\lambda}} = ()\ M$, and then drawing upon the soundness of $\mathbf{Reg}_{=}^{+}$ with respect to equality of strongly regular λ-terms.

We start by introducing a variant of the proof system $\mathbf{Reg}_{\mathbf{0}}^{+}$ in which the formulas are closed, prefixed, $\lambda_{\mathsf{letrec}}$-term-annotated λ-terms.

Definition 1.8.2 (the proof system $\mathbf{ann-Reg}_{\mathbf{0}}^{+}$)**.** The formulas of the proof system $\mathbf{ann-Reg}_{\mathbf{0}}^{+}$ are closed expressions of the form $(\lambda\vec{x})\ L : M$ with $\vec{x}$ a variable prefix vector, $\lambda\vec{x}.\,L$ a $\lambda_{\mathsf{letrec}}$-term, and $\lambda\vec{x}.\,M$ a λ-term. The axioms and rules of $\mathbf{ann-Reg}_{\mathbf{0}}^{+}$ are annotated versions of the axioms and rules of the proof system $\mathbf{Reg}_{\mathbf{0}}^{+}$ from definition 1.6.4 and fig. 1.15, and are displayed in fig. 1.23.

$$\frac{}{(\lambda\vec{x}y)\ y : y}\ \mathsf{0} \qquad \frac{(\lambda\vec{x}y)\ L : M}{(\lambda\vec{x})\ \lambda y.\, L : \lambda y.\, M}\ \lambda$$

$$\frac{(\lambda\vec{x})\ L_0 : M_0 \qquad (\lambda\vec{x})\ L_1 : M_1}{(\lambda\vec{x})\ L_0\ L_1 : M_0\ M_1}\ @$$

$$\frac{(\lambda x_1 \ldots x_{n-1})\ L : M}{(\lambda x_1 \ldots x_n)\ L : M}\ \mathsf{S} \quad \text{(if the binding } \lambda x_n \text{ is vacuous)}$$

$$\frac{\begin{array}{c}[(\lambda\vec{x})\ \mathsf{c}_u : M]^u \\ \mathcal{D}_0 \\ (\lambda\vec{x})\ L[u := \mathsf{c}_u] : M\end{array}}{(\lambda\vec{x})\ (\mathsf{let}\ u = L\ \mathsf{in}\ u) : M}\ \mathrm{FIX}, u \quad \begin{array}{l}\text{(if } |\mathcal{D}_0| \geq 1, \text{ and } |\vec{y}| \geq |\vec{x}| \text{ for all}\\ (\lambda\vec{y})\ N \text{ on threads in } \mathcal{D}_0 \text{ from open}\\ \text{assumptions } ((\lambda\vec{x})\ u : M)^u \text{ down)}\end{array}$$

Figure 1.23. Annotated natural-deduction style proof system $\mathbf{ann-Reg_0^+}$ for strongly regular λ-terms, a version of $\mathbf{Reg_0^+}$ with $\lambda_{\mathsf{letrec}}$-terms as annotations.

Remark 1.8.3. For an example that illustrates why we have chosen to formulate an annotated version only of the proof system $\mathbf{Reg_0^+}$, but not of $\mathbf{Reg^+}$, please see example 1.8.6.

The following proposition is a statement that is entirely analogous to proposition 1.6.7.

Proposition 1.8.4 (cycles are guarded)**.** For all for all instances ι of the rule FIX in a derivation $\mathcal{D}$ (possibly with open assumptions) in $\mathbf{ann-Reg_0^+}$ it holds: every thread from ι upwards to a marked assumption that is discharged at ι passes at least one instance of a rule (λ) or $(@)$.

The lemma below states a straightforward connection between derivations in $\mathbf{Reg_0^+}$ and derivations in its annotated version $\mathbf{ann-Reg_0^+}$.

Lemma 1.8.5 (from $\mathbf{Reg_0^+}$- to $\mathbf{ann-Reg_0^+}$-derivations, and back)**.** The following transformations are possible between derivations in $\mathbf{Reg_0^+}$ and derivations in $\mathbf{ann-Reg_0^+}$:

(i) Every derivation $\mathcal{D}$ in $\mathbf{Reg_0^+}$ with conclusion $(\lambda\vec{x})\ M$ can be transformed into a derivation $\hat{\mathcal{D}}$ in $\mathbf{ann-Reg_0^+}$ with conclusion $(\lambda\vec{x})\ L : M$ such

that there is a bijective correspondence between marked assumptions $((\lambda\vec{y})\ M)^u$ in $\mathcal{D}$ and marked assumptions $((\lambda\vec{y})\ u : M)^u$ in $\hat{\mathcal{D}}$. (As a consequence, $\hat{\mathcal{D}}$ is a closed derivation if $\mathcal{D}$ is closed.) More precisely, $\hat{\mathcal{D}}$ can be obtained from $\mathcal{D}$ by replacing every term occurrence $(\lambda\vec{y})\ N$ by an occurrence of $(\lambda\vec{y})\ P : N$ for a prefixed $\lambda_{\mathsf{letrec}}$-term $(\lambda\vec{y})\ P$ with the property that every prefix variable y_i bound in P is also bound in N. Thereby occurrences of marked assumptions and axioms $\mathsf{0}$ in $\mathcal{D}$ give rise to occurrences of marked assumptions and axioms $\mathsf{0}$ in $\hat{\mathcal{D}}$, respectively; instances of the $\mathbf{Reg}_\mathbf{0}^+$-rules λ, @, S, and FIX in $\mathcal{D}$ give rise to instances of $\mathbf{ann-Reg}_\mathbf{0}^+$-rules λ, @, S, and FIX in $\hat{\mathcal{D}}$, respectively.

(ii) From every closed derivation $\mathcal{D}$ in $\mathbf{ann-Reg}_\mathbf{0}^+$ with conclusion $(\lambda\vec{x})\ L : M$ a closed derivation $\check{\mathcal{D}}$ in $\mathbf{Reg}_\mathbf{0}^+$ with conclusion $(\lambda\vec{x})\ M$ can be obtained by dropping the annotations with $\lambda_{\mathsf{letrec}}$-terms.

Proof. Statement (i) of the lemma can be established through a proof by induction on the depth $|\mathcal{D}|$ of a derivation $\mathcal{D}$ in $\mathbf{Reg}_\mathbf{0}^+$ with possibly open assumptions. In the base case, axioms $(\mathsf{0})$ of $\mathbf{Reg}_\mathbf{0}^+$ are annotated to axioms $(\mathsf{0})$ of $\mathbf{ann-Reg}_\mathbf{0}^+$, and marked assumptions $((\lambda\vec{y})\ N)^u$ in $\mathbf{Reg}_\mathbf{0}^+$ to marked assumptions $((\lambda\vec{y})\ \mathsf{c}_u : N)^u$. In the induction step it has to be shown that a derivation $\mathcal{D}$ in $\mathbf{Reg}_\mathbf{0}^+$ with immediate subderivation $\mathcal{D}_0$ can be annotated to a derivation $\hat{\mathcal{D}}$ in $\mathbf{ann-Reg}_\mathbf{0}^+$, using the induction hypothesis which guarantees that an annotated version $\hat{\mathcal{D}}_0$ of $\mathcal{D}_0$ has already been obtained. Then for obtaining $\hat{\mathcal{D}}$ from $\hat{\mathcal{D}}_0$ the fact is used that the rules in $\mathbf{ann-Reg}_\mathbf{0}^+$ uniquely determine the annotation in the conclusion of an instance once the annotation(s) in the premise(s) (and in the case of FIX additionally the annotation markers used in the assumptions that are discharged) are given. In order to establish that instances of S in $\mathcal{D}$ give rise to corresponding instances of S in $\hat{\mathcal{D}}$, the part of the induction hypothesis is used which guarantees that the $\lambda_{\mathsf{letrec}}$-term annotation in the premise contains not more variable bindings than the λ-term it annotates.

Statement (ii) of the lemma is a consequence of the fact that, by dropping the $\lambda_{\mathsf{letrec}}$-term-annotations, every instance of a rule of $\mathbf{ann-Reg}_\mathbf{0}^+$ give rise to an instance of the corresponding rule in $\mathbf{Reg}_\mathbf{0}^+$. Formally the statement can again be established by induction on the depth of derivations in $\mathbf{ann-Reg}_\mathbf{0}^+$. □

Example 1.8.6. The derivation $\mathcal{D}_l$ in $\mathbf{Reg}_\mathbf{0}^+$ from example 1.6.8 on the left can be annotated, as described by lemma 1.8.5 (ii), to obtain the following derivation $\hat{\mathcal{D}}_l$ in $\mathbf{ann-Reg}_\mathbf{0}^+$:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{(()\ \mathsf{c}_u : M)^u}{(\lambda x)\ \mathsf{c}_u : M}\,\mathsf{S}}{(\lambda xy)\ \mathsf{c}_u : M}\,\mathsf{S} \quad \dfrac{}{(\lambda xy)\ y : y}\,\mathsf{0}}{(\lambda xy)\ \mathsf{c}_u\ y : M\ y}\,@ \quad \dfrac{\dfrac{}{(\lambda x)\ x : x}\,\mathsf{0}}{(\lambda xy)\ x : x}\,\mathsf{S}}{(\lambda xy)\ \mathsf{c}_u\ y\ x : M\ y\ x}\,@}{(\lambda x)\ \lambda y.\, \mathsf{c}_u\ y\ x : M\ y\ x}\,\lambda}{()\ \lambda xy.\, \mathsf{c}_u\ y\ x : \lambda xy.\, M\ y\ x}\,\lambda}{()\ (\mathsf{let}\ u = \lambda xy.\, u\ y\ x\ \mathsf{in}\ f) : M}\,\mathrm{FIX}, u$$

Note that the term in the conclusion, which has been extracted by the annotation procedure, is actually the same as the $\lambda_{\mathsf{letrec}}$-term $\mathsf{let}\ f = \lambda xy.\, f\ y\ x\ \mathsf{in}\ f$ which was used in example 1.6.8 to define M as its infinite unfolding.

Furthermore note that, in a variant of $\mathbf{ann-Reg_0^+}$ in which the '$\mathbf{Reg_0^+}$-addition' (concerning abstraction prefix lengths) to the side-condition of FIX is dropped, the derivation $\mathcal{D}_r$ in example 1.6.8 on the right could be annotated to obtain the following proof tree $\hat{\mathcal{D}}_r$:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{((\lambda x)\ \mathsf{c}_u : \lambda y.\, M\ y\ x)^u}{()\ \lambda x.\, \mathsf{c}_u : M}\,\lambda}{(\lambda x)\ \lambda x.\, \mathsf{c}_u : M}\,\mathsf{S}}{(\lambda xy)\ \lambda x.\, \mathsf{c}_u : M}\,\mathsf{S} \quad \dfrac{}{(\lambda xy)\ y : y}\,\mathsf{0}}{(\lambda xy)\ (\lambda x.\, \mathsf{c}_u)\ y : M\ y}\,@ \quad \dfrac{\dfrac{}{(\lambda x)\ x : x}\,\mathsf{0}}{(\lambda xy)\ x : x}\,\mathsf{S}}{(\lambda xy)\ (\lambda x.\, \mathsf{c}_u)\ y\ x : M\ y\ x}\,@}{(\lambda x)\ \lambda y.\, (\lambda x.\, \mathsf{c}_u)\ y\ x : \lambda y.\, M\ y\ x}\,\lambda}{(\lambda x)\ \mathsf{let}\ u = \lambda y.\, u\ y\ x\ \mathsf{in}\ u : \lambda y.\, M\ y\ x}\,\iota, u}{()\ \lambda x.\, \mathsf{let}\ u = \lambda y.\, u\ y\ x\ \mathsf{in}\ u : M}\,\lambda$$

Observe that, equally as was the case for $\mathcal{D}_r$, also in $\hat{\mathcal{D}}_r$ there occurs, on the thread between the marked assumption at the top and the rule instance ι at which this assumption is discharged, a formula, namely $()\ u : M$, that has a shorter abstraction prefix than the formula in the premise and conclusion of ι as well as in the assumption. Thus ι is not an instance of the rule FIX in $\mathbf{ann-Reg_0^+}$.

Furthermore note that the $\lambda_{\mathsf{letrec}}$-term extracted by $\hat{\mathcal{D}}_r$ does not unfold to M, and hence does not express M. This example shows that the side-condition on instances of FIX in $\mathbf{ann-Reg_0^+}$ cannot be weakened to the form used for

$$\frac{\{[(\lambda\vec{x})\ \mathsf{c}_{f_i} : M_i]\}_{i=1,\ldots,n} \quad \mathcal{D}_j \quad \{\ldots(\lambda\vec{x})\ L_j[\vec{f} := \vec{\mathsf{c}}_{\vec{f}}] : M_j \ldots\}_{j=0,\ldots,n}}{(\lambda\vec{x})\ (\mathsf{let}\ f_1 = L_1 \ldots f_n = L_n\ \mathsf{in}\ L_0) : M_0}\ \mathrm{FIX}_{\mathsf{letrec}}$$

where $\mathsf{c}_{f_1}, \ldots, \mathsf{c}_{f_n}$ are distinct constants fresh for $L_0, \ldots, L_n$, and substitutions $L_l[\vec{f} := \vec{\mathsf{c}}_{\vec{f}}]$ stands for $L_l[f_1 := \mathsf{c}_{f_1}, \ldots, f_n := \mathsf{c}_{f_n}]$.

side-conditions: $|\vec{y}| \geq |\vec{x}|$ holds for the prefix length of every $(\lambda\vec{y})\ N$ on a thread in $\mathcal{D}_j$ for $0 \leq j \leq n$ from an open assumptions $((\lambda\vec{x})\ \mathsf{c}_{f_i})^{u_i}$ downwards; for bottommost instances: the arising derivation is guarded on access path cycles.

Figure 1.24. The proof system $\mathbf{ann-Reg^{+}_{letrec}}$ for $\lambda_{\mathsf{letrec}}$-terms arises from the proof system $\mathbf{ann-Reg^{+}_{0}}$ (see fig. 1.23) by replacing the rule FIX with the rule $\mathrm{FIX}_{\mathsf{letrec}}$.

the rule FIX in $\mathbf{Reg^{+}}$ when the aim is to extract a $\lambda_{\mathsf{letrec}}$-term that unfolds to the infinite λ-term in the conclusion.

The central property of the proof system $\mathbf{ann-Reg^{+}_{0}}$ still remains to be shown: that the $\lambda_{\mathsf{letrec}}$-terms in the conclusion of a derivation in this system does actually unfold to the infinite λ-term in the conclusion. This will be established below in lemma 1.8.11 and theorem 1.8.12. But as an intermediary proof system that will allow us to use results about the proof system $\mathbf{Reg^{+}_{letrec}}$ from section 1.6, we also introduce an annotated version of the rule letrec in $\mathbf{Reg^{+}_{letrec}}$, and an according annotated proof system.

Definition 1.8.7 (the proof system $\mathbf{ann-Reg^{+}_{letrec}}$)**.** The proof system $\mathbf{ann-Reg^{+}_{letrec}}$ arises from $\mathbf{Reg^{+}_{0}}$ by replacing the rule FIX by the rule $\mathrm{FIX}_{\mathsf{letrec}}$ in fig. 1.24, an annotated version of the rule $\mathrm{FIX}_{\mathsf{letrec}}$ from definition 1.6.20 and fig. 1.21. The side-condition on bottommost instances of $\mathrm{FIX}_{\mathsf{letrec}}$ to be guarded on access path cycles is analogous as explained in definition 1.6.20.

Proposition 1.8.8 (from $\mathbf{ann-Reg^{+}_{letrec}}$- to $\mathbf{Reg^{+}_{letrec}}$-derivations)**.** Let $\mathcal{D}$ be a closed derivation in $\mathbf{ann-Reg^{+}_{letrec}}$ with conclusion $()\ L : M$. Then a closed

derivation $\check{\mathcal{D}}$ in $\mathbf{Reg}^{+}_{\mathsf{letrec}}$ with conclusion $()\ L$ can be obtained by removing the λ-terms in $\mathcal{D}$ while keeping the $\lambda_{\mathsf{letrec}}$-term-annotations.

Proposition 1.8.9 (from $\mathbf{ann-Reg}^{+}_{\mathbf{0}}$- to $\mathbf{ann-Reg}^{+}_{\mathsf{letrec}}$-derivations)**.** Every derivation $\mathcal{D}$ in $\mathbf{ann-Reg}^{+}_{\mathbf{0}}$ can be transformed into a derivation $\mathcal{D}'$ in $\mathbf{ann-Reg}^{+}_{\mathsf{letrec}}$ with the same conclusion and with the same open assumption classes.

Proof. First note that $\mathbf{ann-Reg}^{+}_{\mathbf{0}}$ and $\mathbf{ann-Reg}^{+}_{\mathsf{letrec}}$ differ only by the specific version of assumption-discharging rule in the system, FIX in $\mathbf{ann-Reg}^{+}_{\mathbf{0}}$ and $\mathrm{FIX}_{\mathsf{letrec}}$ in $\mathbf{ann-Reg}^{+}_{\mathsf{letrec}}$. For showing the proposition, let $\mathcal{D}$ be a derivation in $\mathbf{ann-Reg}^{+}_{\mathbf{0}}$.

We define a proof tree $\mathcal{D}'$, (intended to be a derivation in $\mathbf{ann-Reg}^{+}_{\mathsf{letrec}}$) by repeatedly replacing topmost occurrences of FIX at the bottom of subderivations of the form as depicted in fig. 1.23, by simulating subderivations of the form:

$$\dfrac{\begin{array}{c}[(\lambda\vec{x})\ \mathsf{c}_u\ :\ M]^u \\ \mathcal{D}_0 \\ ((\lambda\vec{x})\ \mathsf{c}_u\ :\ M)^u \qquad (\lambda\vec{x})\ L[u := \mathsf{c}_u]\ :\ M\end{array}}{(\lambda\vec{x})\ (\mathsf{let}\ u = L\ \mathsf{in}\ u)\ :\ M}\ \mathrm{FIX}_{\mathsf{letrec}}, u$$

until all occurrences of instances of FIX have been replaced by instances of $\mathrm{FIX}_{\mathsf{letrec}}$. The result is a proof tree with axioms and rules of $\mathbf{ann-Reg}^{+}_{\mathsf{letrec}}$ + $\mathrm{FIX}^{-}_{\mathsf{letrec}}$, with the same conclusion and the same classes of open assumptions as $\mathcal{D}$, but in which rule instances carrying the label $\mathrm{FIX}_{\mathsf{letrec}}$ might actually be instances of $\mathrm{FIX}^{-}_{\mathsf{letrec}}$, unless actually proven (as will be done below) to be instances of $\mathrm{FIX}_{\mathsf{letrec}}$.

Now first note that, due to the form of the introduced instances of $\mathrm{FIX}_{\mathsf{letrec}}$, every formula occurrence in $\mathcal{D}$ is reachable on an access path of $\mathcal{D}'$. Second, note that relative access paths π' in $\mathcal{D}'$ starting at the conclusion of an instance ι' of $\mathrm{FIX}_{\mathsf{letrec}}$ up to a marked assumption that is discharged at ι' descend from a thread π in $\mathcal{D}$ from the conclusion of an application ι of FIX up to a marked assumption that is discharged at ι. Since by proposition 1.8.4 the thread π' passes at least one instance of a rule (λ) or $(@)$, this is also the case for π. As a consequence, all cycles on relative access paths are guarded. Thus $\mathcal{D}$ is guarded. Hence all occurrences of rule names $\mathrm{FIX}_{\mathsf{letrec}}$ in $\mathcal{D}'$ rightly label occurrences of this rule, and $\mathcal{D}'$ is a derivation in $\mathbf{ann-Reg}^{+}_{\mathsf{letrec}}$, which moreover is guarded. □

Example 1.8.10. The closed derivation $\hat{\mathcal{D}}_l$ in example 1.8.6 can be transformed into the following closed derivation in $\mathbf{ann-Reg}^{+}_{\mathsf{letrec}}$:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{(()\ \mathsf{c}_u : M)^u}{(\lambda x)\ \mathsf{c}_u : M}\,\mathsf{S} \quad \dfrac{}{(\lambda x)\ x : x}\,\mathsf{0}}{(\lambda x)\ \mathsf{c}_u\ x : M\ x}\,@ \quad \dfrac{}{(\lambda xy)\ y : y}\,\mathsf{0}}{(\lambda xy)\ \mathsf{c}_u\ x : M\ x}\,\mathsf{S} \quad \dfrac{}{(\lambda xy)\ y : y}\,\mathsf{0}}{(\lambda xy)\ \mathsf{c}_u\ x\ y : M\ x\ y}\,@}{(\lambda x)\ \lambda y.\, \mathsf{c}_u\ x\ y : M\ x\ y}\,\lambda}{() \ \lambda xy.\, \mathsf{c}_u\ x\ y : \lambda xy.\, M\ x\ y}\,\lambda \qquad (()\ \mathsf{c}_u : M)^u}{()\ (\mathsf{let}\ u = \lambda xy.\, u\ x\ y\ \mathsf{in}\ u) : M}\,\mathrm{FIX}_{\mathsf{letrec}}, u$$

Now we concentrate on the remaining matter of proving that the $\lambda_{\mathsf{letrec}}$-term obtained by the annotation process from a closed derivation in $\mathbf{Reg}^{+}_{\mathbf{0}}$ to one in $\mathbf{ann-Reg}^{+}_{\mathbf{0}}$ does indeed unfold to the λ-term it annotates. For this, we establish a proof-theoretic transformation from derivations in $\mathbf{ann-Reg}^{+}_{\mathbf{0}}$ to derivations in $\mathbf{Reg}^{+}_{=}$.

Lemma 1.8.11 (from $\mathbf{ann-Reg}^{+}_{\mathbf{0}}$- to $\mathbf{Reg}^{+}_{=}$-derivations)**.** Let $\mathcal{D}$ be a closed derivation in $\mathbf{ann-Reg}^{+}_{\mathbf{0}}$ with conclusion $()\ L : M$. Then $[\![L]\!]_{\boldsymbol{\lambda}}{\downarrow}$, and $\mathcal{D}$ can be transformed into a closed derivation $\mathcal{D}'$ in $\mathbf{Reg}^{+}_{=}$ with conclusion $()\ [\![L]\!]_{\boldsymbol{\lambda}} = ()\ M$ by:

- replacing each formula occurrence o of $(\lambda\vec{y})\ P : N$ in $\mathcal{D}$ by an occurrence of the formula $[\![(\lambda\vec{y})\ \mathsf{let}\ B\ \mathsf{in}\ \tilde{P}]\!]_{\boldsymbol{\lambda}} = (\lambda\vec{y})\ N$ in $\mathcal{D}'$, where B arises as the union of all outermost binding groups in conclusions of instances of FIX at or below o, and where $(\lambda\vec{y})\ P = (\lambda\vec{y})\ \tilde{P}[\vec{f} := \vec{\mathsf{c}}_{\vec{f}}]$ and $\vec{f}$ is comprised of the function variables occurring in B and $\vec{\mathsf{c}}_{\vec{f}}$ are distinct constants for $\vec{f}$ as chosen by $\mathcal{D}$; the unfoldings involved here are always defined.

Proof. Let $\mathcal{D}$ be a closed derivation in $\mathbf{ann-Reg}^{+}_{\mathbf{0}}$ with conclusion $()\ L : M$.

By proposition 1.8.9, $\mathcal{D}$ can be transformed into a closed derivation $\mathcal{D}_1$ in $\mathbf{ann-Reg}^{+}_{\mathsf{letrec}}$ with the same conclusion. Due to proposition 1.8.8, by dropping the infinite terms in $\mathcal{D}_1$, a derivation $\mathcal{D}_2$ in $\mathbf{Reg}^{+}_{\mathsf{letrec}}$ with conclusion $()\ L$ can be obtained. Then it follows from theorem 1.6.26 that $[\![L]\!]_{\boldsymbol{\lambda}}{\downarrow}$, that is, that L unfolds to a λ-term.

We have to show that the transformation of $\mathcal{D}$ into $\mathcal{D}'$ as described in the statement of the lemma is, on the one hand, possible (that is, the unfolding of each prefixed $\lambda_{\mathsf{letrec}}$-term is indeed defined), and on the other hand, that the

proof tree $\mathcal{D}'$ obtained by these replacements is indeed a valid derivation in $\mathbf{Reg}^{+}_{=}$.

We argue for the possibility of these replacements and for their correctness locally, that is by carrying out the replacements from the bottom of $\mathcal{D}$ upwards, thereby recognising for every replacement step that it is possible, and that it indeed produces a valid inference in $\mathbf{Reg}^{+}_{\mathbf{0}}$.

As a typical example of the arguments necessary to establish this fact, we consider a derivation $\mathcal{D}$ in $\mathbf{ann}-\mathbf{Reg}^{+}_{\mathbf{0}}$ with in it an instance of (λ) that immediately succeeds an instance of FIX:

$$\dfrac{\dfrac{\begin{array}{c}[(\lambda\vec{x}y)\ \mathsf{c}_u\ :\ M_0]^u\\ \mathcal{D}_0\\ (\lambda\vec{x}y)\ L_0[u := \mathsf{c}_u]\ :\ M_0\end{array}}{(\lambda\vec{x}y)\ \mathsf{let}\ u = L_0\ \mathsf{in}\ u\ :\ M_0}\ \mathrm{FIX}, u}{(\lambda\vec{x})\ (\lambda y.\,\mathsf{let}\ u = L_0\ \mathsf{in}\ u)\ :\ \lambda y.\,M_0}\ \lambda$$
$$\mathcal{D}'_{00}$$
$$()\ L\ :\ M$$

According to the statement of the lemma, $\mathcal{D}$ is transformed into the following $\mathbf{Reg}^{+}_{=}$-proof-tree:

$$\dfrac{\dfrac{\begin{array}{c}\ldots\ (\llbracket(\lambda\vec{x}y)\ \mathsf{let}\ B_0,\ u = \tilde{L}'_0,\ B'\ \mathsf{in}\ u\rrbracket_{\boldsymbol{\lambda}} = (\lambda\vec{x}y)\ M_0)^u\ \ldots\\ \mathcal{D}'_0\\ \llbracket(\lambda\vec{x}y)\ \mathsf{let}\ B_0,\ u = \tilde{L}'_0\ \mathsf{in}\ \tilde{L}'_0\rrbracket_{\boldsymbol{\lambda}} = (\lambda\vec{x}y)\ M_0\end{array}}{\llbracket(\lambda\vec{x}y)\ \mathsf{let}\ B_0,\ u = \tilde{L}'_0\ \mathsf{in}\ u\rrbracket_{\boldsymbol{\lambda}} = (\lambda\vec{x}y)\ M_0}\ \mathrm{FIX}, u}{\llbracket(\lambda\vec{x})\ \mathsf{let}\ B_0\ \mathsf{in}\ \lambda y.\,\mathsf{let}\ u = \tilde{L}'_0\ \mathsf{in}\ u\rrbracket_{\boldsymbol{\lambda}} = (\lambda\vec{x})\ \lambda y.\,M_0}\ \lambda$$
$$\mathcal{D}'_{00}$$
$$\llbracket()\ \mathsf{let}\ \ \mathsf{in}\ L\rrbracket_{\boldsymbol{\lambda}} = ()\ M$$

where B_0 arises as the union of all outermost binding groups in conclusions of instances of FIX strictly below the visible instance of FIX, and B' is the union of all outermost binding groups in conclusions of instances of FIX strictly above the visible instance of FIX and below the indicated marked assumptions (this binding group differs for different marked assumptions of this assumption class), and where $\tilde{L}'_0$ is the result of replacing in L_0 all occurrences of constants c_f by the function variable f from which it originates.

Now assuming that the unfolding in the conclusion of the visible instance of (@) has been shown to exist, we want to recognise that this instance and the instance of FIX above are valid instances in $\mathbf{Reg_0^+}$. For the instance of (λ) we have to show:

$$\begin{aligned}&[\![(\lambda\vec{x})\ \mathsf{let}\ B_0\ \mathsf{in}\ \lambda y.\,\mathsf{let}\ u = \tilde{L}_0'\ \mathsf{in}\ u]\!]_{\boldsymbol{\lambda}}{\downarrow}\\ &\Longrightarrow \exists\lambda\vec{x}y.\,N_0\ \ [\![(\lambda\vec{x}y)\ \mathsf{let}\ B_0,\, u = \tilde{L}_0'\ \mathsf{in}\ \tilde{L}_0']\!]_{\boldsymbol{\lambda}}{\downarrow} = (\lambda\vec{x}y)\ N_0 \wedge\\ &\qquad\qquad [\![(\lambda\vec{x})\ \mathsf{let}\ B_0\ \mathsf{in}\ \lambda y.\,\mathsf{let}\ u = \tilde{L}_0'\ \mathsf{in}\ u]\!]_{\boldsymbol{\lambda}} = (\lambda\vec{x})\ \lambda y.\,N_0\end{aligned}$$

This, however, is a straightforward consequence of the following $\to_{\triangledown}$-rewrite-steps:

$$\begin{aligned}(\lambda\vec{x})\ \mathsf{let}\ B_0\ \mathsf{in}\ \lambda y.\,\mathsf{let}\ u = \tilde{L}_0'\ \mathsf{in}\ u &\to_{\triangledown.\lambda} \quad (\lambda\vec{x})\ \lambda y.\,\mathsf{let}\ B_0\ \mathsf{in}\ \mathsf{let}\ u = \tilde{L}_0'\ \mathsf{in}\ u\\ &\to_{\triangledown.\mathsf{letrec}} \quad (\lambda\vec{x})\ \lambda y.\,\mathsf{let}\ B_0, u = \tilde{L}_0'\ \mathsf{in}\ u\end{aligned}$$

in view of the fact that, by lemma 0.7.7, unfolding is uniquely normalising (in at most ω steps). And for the instance of FIX we have to show:

$$\underbrace{[\![(\lambda\vec{x}y)\ \mathsf{let}\ B_0,\, u = \tilde{L}_0'\ \mathsf{in}\ u]\!]_{\boldsymbol{\lambda}}{\downarrow}}_{=(\star)} \Longrightarrow [\![(\lambda\vec{x}y)\ \mathsf{let}\ B_0,\, u = \tilde{L}_0'\ \mathsf{in}\ \tilde{L}_0']\!]_{\boldsymbol{\lambda}}{\downarrow} = (\star)$$
$$\wedge\ [\![(\lambda\vec{x}y)\ \mathsf{let}\ B_0,\, u = \tilde{L}_0',\, B'\ \mathsf{in}\ u]\!]_{\boldsymbol{\lambda}}{\downarrow} = (\star)$$

(actually the statement as in the second line has to be shown for every binding-group B' that occurs for marked assumptions discharged at the instance of FIX). This implication is a consequence of the $\to_{\triangledown}$-rewrite-steps:

$$\begin{aligned}(\lambda\vec{x}y)\ \mathsf{let}\ B_0,\, u = \tilde{L}_0'\ \mathsf{in}\ u &\to_{\triangledown.\mathrm{rec}} (\lambda\vec{x}y)\ \mathsf{let}\ B_0,\, u = \tilde{L}_0'\ \mathsf{in}\ \tilde{L}_0'\\ (\lambda\vec{x}y)\ \mathsf{let}\ B_0,\, u = \tilde{L}_0',\, B'\ \mathsf{in}\ u &\to_{\triangledown.\mathrm{red}} (\lambda\vec{x}y)\ \mathsf{let}\ B_0,\, u = \tilde{L}_0'\ \mathsf{in}\ u\end{aligned}$$

again in view of the statement of lemma 0.7.7.

The arguments used here are typical, and can be carried out similarly also for showing that axioms (0), and instances of rules (@) and (S) in $\mathbf{ann-Reg_0^+}$-derivations give rise to, under the transformation described in the statement of the lemma, valid instances of axioms (0), and instances of (@) and (S), respectively, in $\mathbf{Reg_=^+}$-derivations. □

Theorem 1.8.12 ($\mathbf{ann-Reg_0^+}$ and unfolding semantics)**.** If $\vdash_{\mathbf{ann-Reg_0^+}} () \ L : M$ holds for a $\lambda_{\mathsf{letrec}}$-term L and a λ-term M, then L unfolds to M.

Proof. Suppose that $\mathcal{D}$ is a closed derivation in $\mathbf{ann-Reg_0^+}$ with conclusion $()\ L : M$. Lemma 1.8.11 entails that L unfolds to a λ-term, and moreover, that $\mathcal{D}$ can be transformed into a closed derivation $\mathcal{D}'$ in $\mathbf{Reg_=^+}$ with conclusion $()\ [\![L]\!]_{\boldsymbol{\lambda}} = ()\ M$. Then it follows by theorem 1.6.19 (applying soundness of $\mathbf{Reg_0^+}$ with respect to the property of $\lambda_{\mathsf{letrec}}$-terms to unfold to a λ-term that $[\![L]\!]_{\boldsymbol{\lambda}} = M$, and hence that $L \to_{\triangledown}^{\omega} M$. In this way we have found a $\lambda_{\mathsf{letrec}}$-term L that expresses M. □

We now arrive at our main characterisation result.

Theorem 1.8.13 ($\lambda_{\mathsf{letrec}}$-expressibility ~ strong regularity)**.** A λ-term is $\lambda_{\mathsf{letrec}}$-expressible if and only if it is strongly regular.

Proof. Let M be a λ-term.

The direction "⇒" is the statement of theorem 1.6.28.

For showing the direction "⇐" in the statement of the theorem, we assume that M is strongly regular. Then by theorem 1.6.15 (ii), there exists a closed derivation $\mathcal{D}$ in $\mathbf{Reg^+}$ with conclusion $()\ M$. Due to lemma 1.8.5 (i), $\mathcal{D}$ can be transformed into a derivation $\hat{\mathcal{D}}$ in $\mathbf{ann-Reg_0^+}$ with conclusion $()\ L : M$, for some $\lambda_{\mathsf{letrec}}$-term L. Then it follows by theorem 1.8.12 that the $\lambda_{\mathsf{letrec}}$-term L expresses M. □

As an immediate consequence of theorem 1.8.13 and of corollary 1.7.22 we obtain the following theorem, a summary of our main results:

Theorem 1.8.14. For all λ-terms M the following statements are equivalent:

(i) M is $\lambda_{\mathsf{letrec}}$-expressible.

(ii) M is strongly regular.

(iii) M is regular, and it only contains finite binding–capturing chains.

1.9 λ-transition-graphs

§ 1.9.1 (Overview)**.** In this section we introduce the concept of λ-transition-graphs. λ-transition-graphs are a nameless graphical representations of λ-terms that arise naturally from the decomposition systems from this chapter. The λ-transition-graph of M is in fact almost identical to M's reduction graph with respect to some scope$^+$-delimiting strategy for $\boldsymbol{Reg^+}$. The only difference is that the former is a (pointed) LTS and the latter an ARS. As explained earlier in

remark 1.3.19 these formalisms are essentially same, except for the definition of bisimulation, which for LTSs is sensitive to the transition labels. The main result of this section is a coinduction principle for λ-terms: two λ-terms are equal if and only if they have bisimilar λ-transition-graphs.

§ 1.9.2 (outlook: λ-term-graphs)**.** The λ-transition-graphs which we study in this section will be further developed in chapter 2 into 'λ-term-graphs', which are first-order term graph and serve as a graphical representation of λ-terms and $\lambda_{\mathsf{letrec}}$-terms.

Definition 1.9.3 (transition system induced by an ARS)**.** Let $\mathcal{A} = \langle O, \Phi, \mathrm{src}, \mathrm{tgt}\rangle$ be an ARS (or a sub-ARS (or a sub-ARS of a labelled version) of an ARS) that is induced by a CRS with rules R (see § 1.3.6). Note that, every step in $\Phi : O \times R \times O$ carries information according to from which rule it stems from. By the *LTS induced by* $\mathcal{A}$ we mean the LTS $\mathcal{L}_{\mathcal{A}} = \langle O, R, \Phi\rangle$ in which the steps in $\mathcal{A}$ according to rule ρ are interpreted as transitions with label ρ.

Definition 1.9.4 (transition graph of a λ-term)**.** Let M be a λ-term with reduction graph $(M \twoheadrightarrow_{\mathcal{A}})$ w.r.t. to an ARS $\mathcal{A}$ and let $\langle O, R, T\rangle$ be the LTS induced by $(M \twoheadrightarrow_{\mathcal{A}})$. We call the labelled transition graph $\mathcal{G}_{\mathcal{A}}(M) = \langle O, R, M, T\rangle$ the *transition graph of* M *w.r.t.* $\mathcal{A}$.

Definition 1.9.5 (λ-transition-graph)**.** We call a labelled transition graph $G = \langle S, L, i, T\rangle$ a *λ-transition-graph* if:

- it is connected
- the labels are $L = \{\lambda, \mathsf{S}, @_0, @_1\}$
- there are no infinite paths in G consisting solely of S-transitions
- every state s belongs to one of the following kinds: λ-states, S-states, and @-states, where
  - a λ-state is the source of precisely one λ-transition, and no other transitions: $\{\langle l, t\rangle \mid s \to_l t\} = \{\langle \lambda, u\rangle\}$ for some $u \in S$.
  - a S-state is the source of precisely one S-transition, and no other transitions: $\{\langle l, t\rangle \mid s \to_l t\} = \{\langle \mathsf{S}, u\rangle\}$ for some $u \in S$.
  - a @-state is the source of precisely one $@_0$-transition and one $@_1$-transition, and no other transitions:
    $\{\langle l, t\rangle \mid s \to_l t\} = \{\langle @_0, f\rangle, \langle @_1, x\rangle\}$ for some $f, x \in S$.

Proposition 1.9.6 (eager scope-closure yields λ-transition-graphs)**.** Let $\mathbb{S}^+$ be a scope$^+$-delimiting strategy for Reg^+. For every term $M \in \mathrm{Ter}^\infty(\boldsymbol{(\lambda)})$ the transition graph $\mathcal{G}_{\mathbb{S}^+}(M)$ is a λ-transition-graph. We call $\mathcal{G}_{\mathbb{S}^+}(M)$ *the λ-transition-graph of M with respect to $\mathbb{S}^+$.*

Proof. In the transition graph $\mathcal{G}_{\mathbb{S}^+}(M)$ there cannot be infinitely many successive S-transitions because in the ARS that induces $\mathcal{G}_{\mathbb{S}^+}(M)$, the rewrite relation $\to_{\mathsf{S}}$ is terminating, due to proposition 1.4.24 (v). □

§ 1.9.7 (λ-transition-graphs of $\lambda_{\mathsf{letrec}}$-terms)**.** Along the lines of proposition 1.9.6 we can also view transition graphs of $\lambda_{\mathsf{letrec}}$-terms as λ-transition-graphs; however, unfolding steps must not be included as transitions (let us call them *silent transitions*). As hinted at in § 1.5.11, here it is important to restrict scope/scope$^+$-delimiting strategies to ones that are deterministic in the application of unfolding rules.

Definition 1.9.8 (LTS with silent steps)**.** Let $\mathcal{L}_{\mathcal{A}} = \langle O, R, \Phi\rangle$ be the LTS induced by ARS $\mathcal{A}$ and let R_0 be a subset of R. Then by *the LTS induced by $\mathcal{A}$ with silent R_0-steps* we mean the LTS $\mathcal{L}_{\mathcal{A},R_0} = \langle O, R, T\rangle$ with

$$T := \left\{ \langle o, \rho, o'\rangle \;\middle|\; \begin{array}{l}\text{if } o \twoheadrightarrow_{R_0} \cdot \to_\rho o' \text{ where } \twoheadrightarrow_{R_0} \text{ are steps w.r.t. rules}\\ \text{in } R_0\text{, and } \to_\rho \text{ is a step w.r.t. a rule } \rho \in R \setminus R_0\end{array} \right\}$$

in which the steps in $\mathcal{A}$ according to rules in R_0 are interpreted as silent transitions.

Definition 1.9.9 (LTG with silent steps)**.** Let o be an object of the ARS $\mathcal{A}$ and $\mathcal{L}_{(o\twoheadrightarrow),R_0} = \langle O, R, T\rangle$ an LTS with silent R_0-steps. We call $\mathcal{G}_{\mathcal{A},R_0}(o) := \langle O, R, o, T\rangle$ the *transition graph of o with silent R_0-steps.*

Proposition 1.9.10. Let $\mathbb{S}^+$ be a scope$^+$-delimiting strategy for Reg^+_{letrec}. For every term $L \in \mathrm{Ter}(\boldsymbol{(\lambda_{\mathsf{letrec}})})$, the transition graph $\mathcal{G}_{\mathbb{S}^+,\boldsymbol{R_\nabla}}(L)$ of L is a λ-transition-graph.

Definition 1.9.11 (λ-transition-graph of a $\lambda_{\mathsf{letrec}}$-term)**.** Let $L \in \mathrm{Ter}(\boldsymbol{(\lambda_{\mathsf{letrec}})})$ be a (prefixed) $\mathbb{S}^+$-productive $\lambda_{\mathsf{letrec}}$-term. For a scope$^+$-delimiting strategy $\mathbb{S}^+$ of Reg^+_{letrec} such that L is $\mathbb{S}^+$-productive, we call the transition graph $\mathcal{G}_{\mathbb{S}^+,\boldsymbol{R_\nabla}}(L)$ *the λ-transition-graph of L with respect to $\mathbb{S}^+$.* We also speak of λ-transition-graphs of terms $L \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$ or $M \in \mathrm{Ter}^\infty(\boldsymbol{\lambda})$ by which we refer to the λ-transition-graphs of $()\ L$ and $()\ M$, respectively.

Theorem 1.9.12 (coinduction principle for $\boldsymbol{\lambda}$)**.** For all λ-terms M and N the following statements are equivalent:

(i) $M = N$.

(ii) $\vdash_{\mathbf{EQ}^{\infty}} M = N$.

(iii) M and N have bisimilar λ-transition-graphs.

Proof. In view of proposition 1.6.17 (ii), the logical equivalence between (i) and (ii), and the fact that (i) $\Rightarrow$ (iii) clearly holds, it suffices to show that (iii) $\Rightarrow$ (ii) holds.

For this, suppose that $\mathcal{G}_{\mathbb{S}_1^+}(M)$ and $\mathcal{G}_{\mathbb{S}_2^+}(N)$ are bisimilar for some scope$^+$-delimiting strategies $\mathbb{S}_1^+$ and $\mathbb{S}_2^+$ for *Reg*. But now bisimilarity of these transition graphs guarantees that a derivation $\mathcal{D}$ in $\mathbf{EQ}^{\infty}$ with conclusion $() M = () N$ can be constructed such that all threads in $\mathcal{D}$ correspond to $\rightarrow_{\mathbb{S}_1^+}$-rewrite-sequences on M and to $\rightarrow_{\mathbb{S}_2^+}$-rewrite-sequences on M, respectively. If the construction process is organised in a depth-fair manner (for example, all non-axiom leafs at depth n are extended by appropriate rule instances, before extensions are carried out at depth greater than n), then in the limit a completed derivation $\mathcal{D}^{\infty}$ with conclusion $() M = () N$ is obtained. This establishes $\vdash_{\mathbf{EQ}^{\infty}} M = N$. □

Conjecture 1.9.13 (coinduction principle for $\boldsymbol{\lambda}_{\mathsf{letrec}}$)**.** For all $L_1, L_2 \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$ it holds that $L_1 = L_2$ if and only if L_1 and L_2 have bisimilar λ-transition-graphs.

Proof sketch. In chapter 3 we develop a coinduction principle for $\boldsymbol{\lambda}_{\mathsf{letrec}}$ (theorem 3.6.17) based on λ-term-graphs. λ-term-graphs are introduced in chapter 2 and are closely related to λ-transition-graphs. The conjecture can in all likelihood be validated by theorem 3.6.17 after relating λ-transition-graphs to λ-term-graphs in a formal manner. □

§ 1.9.14 (only ***Reg*$^+$** defines nameless representations)**.** The coinduction principle above holds for transition graphs derived from terms w.r.t. an eager scope$^+$-delimiting strategy for *Reg*$^+$. It cannot be extended to *Reg*, which is witnessed by the following counterexample.

Consider the transition graph $\mathcal{G}_{\mathbb{S}_{\mathrm{eag}}}(M)$ of the term $M = \lambda xy.\, x\ x\ y$ w.r.t. the eager scope-delimiting strategy $\mathbb{S}_{\mathrm{eag}}$ for *Reg*. The corresponding *Reg*-reduction graph is depicted on the left in fig. 1.9. Each of the following terms yields the exact same transition graph w.r.t. to an appropriately-chosen scope-delimiting

strategy for *Reg*. For the two terms in the middle the eager scope-delimiting strategy can be chosen.

$$\lambda xy.\, x\ x\ x \qquad \lambda xy.\, x\ x\ y \qquad \lambda xy.\, y\ y\ x \qquad \lambda xy.\, y\ y\ y$$

§ 1.9.15 (readback for λ-transition-graphs)**.** The understanding of λ-transition-graphs as nameless representations of λ-terms implies that from a such graphs the corresponding λ-term can be extracted. We define a function for this purpose by means of a CRS which implements the assembly of a λ-term from the infinite unfolding of a λ-transition-graph. The function is closely related to the ***Parse***$^+$ in the sense that ***Parse***$^+$ does both destruct and reconstruct the scrutinised term while *rb* only implements the reconstruction.

Definition 1.9.16 (readback for λ-transition-graphs)**.**

$$\begin{aligned} rb : \mathrm{Ter}^\infty(\{\mathsf{0}, \lambda, @, \mathsf{S}\}) &\nrightarrow \mathrm{Ter}^\infty(\boldsymbol{\lambda}) \\ t \mapsto rb(t) :=\ & \text{infinite normal form of } \mathsf{rw}_0(t) \\ & \text{w.r.t. the following CRS:} \end{aligned}$$

$$\begin{aligned} \mathsf{rw}_n(X_1, \ldots, X_n, \lambda(t_0)) &\to \mathsf{abs}([x]\, \mathsf{rw}_{n+1}(X_1, \ldots, X_n, x, t_0)) \\ \mathsf{rw}_n\big(\vec{X}, @(t_0)t_1\big) &\to \mathsf{app}\big(\mathsf{rw}_n\big(\vec{X}, t_0\big), \mathsf{rw}_n\big(\vec{X}, t_1\big)\big) \\ \mathsf{rw}_{n+1}\big(\vec{X}, x, \mathsf{S}(t_0)\big) &\to \mathsf{rw}_n\big(\vec{X}, t_0\big) \\ \mathsf{rw}_n(X_1, \ldots, X_n, \mathsf{0}) &\to X_n \end{aligned}$$

The function is partial because rw_n is unproductive for infinite S-chains. That restriction comes forth accordingly in the definition of λ-transition-graphs (definition 1.9.5). The function is thus complete on the subset of $\mathrm{Ter}^\infty(\{\mathsf{0}, \lambda, @, \mathsf{S}\})$ that is obtained from unfolding a λ-transition-graph.

1.10 Summary

§ 1.10.1 (A CRS for decomposing λ-terms)**.** To characterise the set of $\lambda_{\mathsf{letrec}}^-$-expressible λ-terms we established a framework of formalisms for 'observing' λ-terms coinductively. First we introduced prefixed λ-terms that enrich λ-terms by an abstraction prefix. On the prefixed terms we defined the CRS Reg^+ in which a rewrite sequence corresponds to a deconstruction of a term along one of its paths. In that sense a prefixed term $(\lambda\vec{x})\ M$ can be understood as a 'suspended decomposition' which has not advanced into subterm M yet. Such a decomposition describes a path through the term by observations of the form $\to_\lambda$, $\to_{@_0}$, $\to_{@_1}$, $\to_{\mathsf{S}}$, where the $\to_{\mathsf{S}}$ delimits the scope$^+$ of an abstraction.

§ 1.10.2 (scope/scope$^+$-delimiting strategies)**.** Since there is some freedom as to where scope$^+$-delimiters can be placed, we defined scope$^+$-delimiting strategies to formalise specific possible choices eliminating that freedom and thereby making the observations deterministic except for the forking into the left or the right subterm of an application. By means of scope$^+$-delimiting strategies we formulated two important concepts: strong regularity and λ-transition graphs.

§ 1.10.3 (strong regularity)**.** The intuitive understanding of strong regularity is the property of a λ-term M that from M every 'sufficiently eager' scope$^+$-delimiting strategy can only generate a finite number of terms. We showed that $\lambda_{\mathsf{letrec}}$-expressibility coincides with strong regularity.

§ 1.10.4 (λ-transition-graphs)**.** Every scope$^+$-delimiting strategy defines for each term a λ-transition-graph which can be viewed as a nameless graphical representation very similar to its term graph in de-Bruijn notation with the difference that S-nodes are not restricted to occur only near leafs but can be shared by variables (see § 1.4.12). We established a coinduction principle for λ-transition-graphs that states that two λ-terms are equal if and only if their λ-transition graphs w.r.t. to a scope$^+$-delimiting strategy are bisimilar. The eager scope$^+$-delimiting strategy yields finite λ-transition-graphs for strongly regular λ-terms.

§ 1.10.5 ($\boldsymbol{\lambda}_{\mathsf{letrec}}$)**.** We then adapted the concepts of the CRS for observing terms, scope$^+$-delimiting strategies, and λ-transition-graphs and applied them to $\boldsymbol{\lambda}_{\mathsf{letrec}}$ proving similar results as for $\boldsymbol{\lambda}$.

§ 1.10.6 (proof systems for strong regularity)**.** We provided a proof system that is sound and complete for the notion of strong regularity and which admits finite proofs for strongly regular λ-terms. We define an annotated version of the proof system which not unlike an attribute-grammar definition implements the extraction of a $\lambda_{\mathsf{letrec}}$-term L from a proof for term M in that system, such that L unfolds to the M. We show that every scope$^+$-delimiting strategy induces a proof and that from a proof a corresponding history-aware strategy can be deduced, which suggests a similar correspondence between λ-transition-graphs and proofs.

Chapter 2

Term Graph Representations for Strongly Regular λ-Terms

2.1 Overview

§ 2.1.0 (teaser)**.** Can't we just do bisimulation on $\lambda_{\mathsf{letrec}}$-terms? See also § 3.1.0.

§ 2.1.1 (subject matter)**.** In this chapter we study various classes of higher-order and first-order term graphs (intended as representations for $\lambda_{\mathsf{letrec}}$-terms). We focus on the relation between 'λ-higher-order term graphs' (λ-ho-term-graphs), which are first-order term graphs endowed with a well-behaved scope function, and their representations as 'λ-term-graphs', which are plain first-order term graphs with scope-delimiter vertices that meet certain scoping requirements. Specifically we tackle the question: Which class of first-order term graphs admits a faithful embedding of λ-ho-term-graphs in the sense that (i) the homomorphism-based sharing-order on λ-ho-term-graphs is preserved and reflected, and (ii) the image of the embedding corresponds closely to a natural class (of λ-term-graphs) that is closed under functional bisimulation?

§ 2.1.2 (motivation)**.** We study these graph formalisms in isolation – that is to say without formally connecting them to $\lambda_{\mathsf{letrec}}$-terms. But we do so with a long term goal in mind: to find a graph formalism that is suitable to adequately represent $\lambda_{\mathsf{letrec}}$-terms. Once we have found such a graph formalism, in chapter 3 we relate it back to $\boldsymbol{\lambda}_{\mathsf{letrec}}$, by which we gain further insights concerning unfolding semantics of $\boldsymbol{\lambda}_{\mathsf{letrec}}$ and also to obtain concrete practical methods to analyse and manipulate $\lambda_{\mathsf{letrec}}$-terms.

§ 2.1.3 (methods and formalisms)**.** The term graph formalims we study arise naturally from the term decomposition systems from the previous chapter. They are closely related to the λ-transition-graphs, and we derive different classes of term graphs from the term decomposition systems by a similar approach. We systematically examine which classes of λ-term-graphs satisfy the properties in § 2.1.1.

§ 2.1.4 (results)**.** We identify a particular class of first-order term graphs with these properties. Term graphs of this class are built not only from application, abstraction, and variable vertices, but also scope-delimiter vertices. They have the characteristic feature that the latter two kinds of vertices have backlinks to the corresponding abstraction. This result puts a handle on the concept of subterm sharing for higher-order term graphs, both theoretically and algorithmically: We obtain an easily implementable method for obtaining the maximally shared form of λ-ho-term-graphs. Also, we open up the possibility to transfer properties from first-order term graphs to λ-ho-term-graphs. In fact we prove in this way that the sharing-order on a set of bisimilar λ-ho-term-graphs forms a complete lattice section 2.10.

§ 2.1.5 (outlook)**.** In chapter 3 we use these insights to develop practical applications w.r.t $\boldsymbol{\lambda}_{\mathsf{letrec}}$:

- an efficient test for whether two $\lambda_{\mathsf{letrec}}$-terms have the same unfolding
- a partial order for the amount of subterm sharing in a $\lambda_{\mathsf{letrec}}$-term leading to
  - a notion of maximal sharing for $\lambda_{\mathsf{letrec}}$
  - an efficient mechanism to compute the maximally shared form of a $\lambda_{\mathsf{letrec}}$-term which generalises common subexpression elimination

2.2 Preliminaries

§ 2.2.1 (term graphs)**.** The graph formalisms that we study in this chapter all based on term graphs (in contrast to transition graphs as in the last chapter), which is a natural choice, considering the three types of expressions (abstraction, application, variable occurrence) that make up λ-terms. Here is some notation and terminology revolving around term graphs.

Definition 2.2.2 (term graph)**.** Let Σ be a signature with arity function $\mathrm{ar} : \Sigma \to \mathbb{N}$. A *term graph over* Σ (or a Σ-*term-graph*) is a tuple $\langle V, \mathrm{lab}, \mathrm{args}, r\rangle$ where V is a set of *vertices*, $\mathrm{lab} : V \to \Sigma$ the *(vertex) label function*, $\mathrm{args} : V \to V^*$ the *argument function* that maps every vertex v to the word $\mathrm{args}(v)$ consisting of the $\mathrm{ar}(\mathrm{lab}(v))$ successor vertices of v (hence it holds that $|\mathrm{args}(v)| = \mathrm{ar}(\mathrm{lab}(v))$), and r, the *root*, is a vertex in V. Note that term graphs may have infinitely many vertices.

Definition 2.2.3 (root connected term graphs)**.** We say that such a term graph is *root connected* if every vertex is reachable from the root by a path that arises by repeatedly going from a vertex to one of its successors. We denote by $\mathrm{TG}(\Sigma)$ and by $\mathrm{TG}^-(\Sigma)$ the class of all root-connected term graphs over Σ, and the class of all term graphs over Σ, respectively. By 'term graphs' we will from now on, always mean root-connected term graphs, except in a few situations in which we explicitly state otherwise.

Definition 2.2.4 (successor relations)**.** Let G be a term graph over signature Σ. As useful notation for picking out the i-th vertex, from among the ordered successors $\mathrm{args}(v)$ of a vertex v of G, we define for each $i \in \mathbb{N}$ the indexed edge relation $\rightarrowtail_i \subseteq V \times V$, and additionally the (not indexed) edge relation $\rightarrowtail \subseteq V \times V$, by stipulating for all $v, v' \in V$:

$$v \rightarrowtail_i v' \;:\Leftrightarrow\; \exists v_0, \ldots, v_n \in V \;\; \mathrm{args}(v) = v_0 \ldots v_n \wedge v' = v_i$$
$$v \rightarrowtail v' \;:\Leftrightarrow\; \exists i \in \mathbb{N} \;\; v \rightarrowtail_i v'$$

We write $v \overset{l}{\rightarrowtail}_i v'$ if $v \rightarrowtail_i v' \wedge \mathrm{lab}(v) = l$ holds for $v, v' \in V$, $i \in \mathbb{N}$, $l \in \Sigma$, to indicate the label at the source of an edge.

Definition 2.2.5 (paths)**.** Let $v_0, \ldots, v_n \in V$. A *path* in G is a tuple $\langle v_0, l_1, v_1, l_2, v_2, l_3, \ldots, l_{n-1}, v_{n-1}, l_n, v_n\rangle$ and $n, l_1, \ldots, l_n \in \mathbb{N}$ such that $v_0 \rightarrowtail_{l_1} v_1 \rightarrowtail_{l_2} v_2 \rightarrowtail_{l_3} \ldots \rightarrowtail_{l_n} v_n$ holds; paths will usually be denoted in the latter form, using indexed edge relations.

Definition 2.2.6 (access paths)**.** An *access path* of a vertex v of a term graph G is a path that starts at the root of G, ends in v, and does not visit any vertex twice. Note that every vertex v has at least one access path: since every vertex in a term graph is reachable from the root (see definition 2.2.3), there is a path π from r to v; then an access path of v can be obtained from π by repeatedly cutting out *cycles*, that is, parts of the path between one and the same vertex.

In the following, let $G_1 = \langle V_1, \mathrm{lab}_1, \mathrm{args}_1, r_1\rangle$, $G_2 = \langle V_2, \mathrm{lab}_2, \mathrm{args}_2, r_2\rangle$ be term graphs over signature Σ.

Definition 2.2.7 (homomorphism, functional bisimulation)**.** A *homomorphism*, also called a *functional bisimulation*, from G_1 to G_2 is a morphism from the structure $\langle V_1, \mathrm{lab}_1, \mathrm{args}_1, r_1\rangle$ to the structure $\langle V_2, \mathrm{lab}_2, \mathrm{args}_2, r_2\rangle$, i.e., a function $h : V_1 \to V_2$ such that, for all $v \in V_1$ it holds:

$$h(r_1) = r_2 \qquad \text{(roots)}$$
$$\mathrm{lab}_1(v) = \mathrm{lab}_2(h(v)) \qquad \text{(labels)}$$
$$h^*(\mathrm{args}_1(v)) = \mathrm{args}_2(h(v)) \qquad \text{(arguments)}$$

where h^* is the homomorphic extension (also: pointwise lifted version) of h to words over V_1, i.e. $h^* : V_1^* \to V_2^*$, $v_1 \dots v_n \mapsto h(v_1) \dots h(v_n)$. In this case we write $G_1 \rightrightarrows_h G_2$, or $G_2 \leftleftarrows_h G_1$. And we write $G_1 \rightrightarrows G_2$, or for that matter $G_2 \leftleftarrows G_1$, if there is a homomorphism from G_1 to G_2.

Let $f \in \Sigma$. An *f-homomorphism* from G_1 to G_2 is a homomorphism h from G_1 to G_2 that 'shares' (i.e. maps to the same vertex) only vertices with the label f, i.e. h has the property that $h(v_1) = h(v_2) \Rightarrow \mathrm{lab}_1(v_1) = \mathrm{lab}_1(v_2) = f$ holds for all $v_1 \neq v_2 \in V_1$. If h is an f-homomorphism from G_1 to G_2, then we write $G_1 \rightrightarrows_h^f G_2$ or $G_2 \leftleftarrows_h^f G_1$, or dropping h, $G_1 \rightrightarrows^f G_2$ or $G_2 \leftleftarrows^f G_1$.

Terminology 2.2.8. The terms 'homomorphism' and 'functional bisimulation' will we used interchangeably throughout this document.

Definition 2.2.9 (isomorphism)**.** An *isomorphism* between G_1 and G_2 is a bijective homomorphism $i : V_1 \to V_2$ from G_1 to G_2 (it follows from the homomorphism conditions definition 2.2.7 that also the inverse function $i^{-1} : V_1 \to V_2$ is a homomorphism). If there is an isomorphism between G_1 and G_2, we write $G_1 \sim G_2$, and say that G_1 and G_2 are *isomorphic*. The relation $\sim$ is an equivalence relation on $\mathrm{TG}(\Sigma)$. For every term graph G over Σ we denote the isomorphism equivalence class $[G]_\sim$ by (the boldface letter) $\boldsymbol{G}$.

Definition 2.2.10 (bisimulation, bisimilarity)**.** For $i \in \{1, 2\}$, let $G_i = \langle V_i, \mathrm{lab}_i, \mathrm{args}_i, r_i\rangle$ be term graphs over signature Σ. A *bisimulation* between G_1 and G_2 is a relation $R \subseteq V_1 \times V_2$ such that the following conditions hold, for all $\langle v, v'\rangle \in R$:

$$\langle r_1, r_2\rangle \in R \qquad \text{(roots)}$$
$$\mathrm{lab}_1(v) = \mathrm{lab}_2(v') \qquad \text{(labels)}$$
$$\langle \mathrm{args}_1(v), \mathrm{args}_2(v')\rangle \in R^* \qquad \text{(arguments)}$$

where the extension $R^* \subseteq {V_1}^* \times {V_2}^*$ of R to a relation between words over V_1 and words over V_2 is defined as:

$$R^* := \left\{ \langle v_1 \dots v_n, w_1 \dots w_n \rangle \;\middle|\; \begin{array}{l} n \in \mathbb{N}, \forall i \in \{1, \dots, n\}\ v_i \in V_1, w_i \in V_2 \\ \text{such that } \langle v_i, w_i \rangle \in R \text{ for all } 1 \le i \le n \end{array} \right\}$$

We write $G_1 \leftrightarrows G_2$ if there is a bisimulation between G_1 and G_2, and say that G_1 and G_2 are *bisimilar*. Bisimilarity $\leftrightarrows$ is an equivalence relation on classes $\mathrm{TG}(\Sigma)$ of term graphs over a signature Σ.

An *f-bisimulation* between G_1 and G_2 is a bisimulation between G_1 and G_2 such that its restriction to vertices with labels different from f is a bijective function. If there is an f-bisimulation between G_1 and G_2 we say that G_1 and G_2 are f-bisimilar and write $\leftrightarrows^f$ to indicate *f-bisimilarity*.

The following proposition is a simple but useful reformulation of the definition of homomorphism (definition 2.2.7).

Proposition 2.2.11 (homomorphism)**.** Let $G_i = \langle V_i, \mathrm{lab}_i, \mathrm{args}_i, r_i \rangle$, for $i \in \{1, 2\}$ be term graphs over signature Σ. Let $h : V_1 \to V_2$ be a function. Then h is a homomorphism from G_1 to G_2, if and only if, for all $v, v_1 \in V_1$, $v_2 \in V_2$, and $k \in \mathbb{N}$, the following four statements hold:

$$h(r_1) = r_2 \qquad \text{(roots)}$$

$$\mathrm{lab}(v) = \mathrm{lab}(h(v)) \qquad \text{(labels)}$$

$$h(v_1) = v_2 \wedge \forall v_1' \in V_1\ v_1 \rightarrowtail_k v_1' \Rightarrow \exists v_2' \in V_2\ v_2 \rightarrowtail_k v_2' \wedge h(v_1') = v_2' \qquad \text{(args-forward)}$$

$$h(v_1) = v_2 \wedge \forall v_2' \in V_2\ v_2 \rightarrowtail_k v_2' \Rightarrow \exists v_1' \in V_1\ v_1 \rightarrowtail_k v_1' \wedge h(v_1') = v_2' \qquad \text{(args-backward)}$$

Proposition 2.2.12 (functional bisimulation and paths)**.** For $i \in \{1, 2\}$ let $G_i = \langle V_i, \mathrm{lab}_i, \mathrm{args}_i, r_i \rangle$ be term graphs over signature Σ. Let $h : V_1 \to V_2$ be a homomorphism from G_1 to G_2. Then the following statements hold:

(i) Every path $\pi : v_0 \rightarrowtail_{k_1} v_1 \rightarrowtail_{k_2} v_2 \rightarrowtail_{k_3} \dots \rightarrowtail_{k_{n-1}} v_{n-1} \rightarrowtail_{k_n} v_n$ in G_1 has an image $h(\pi)$ under h in G_2 in the sense that $h(\pi) : h(v_0) \rightarrowtail_{k_1} h(v_1) \rightarrowtail_{k_2} h(v_2) \rightarrowtail_{k_3} \dots \rightarrowtail_{k_{n-1}} h(v_{n-1}) \rightarrowtail_{k_n} h(v_n)$.

(ii) For every $v_0 \in V_0$ and $v_0' \in V_1$ with $h(v_0) = v_0'$ it holds that every path $\pi' : v_0' \rightarrowtail_{k_1} v_1' \rightarrowtail_{k_1} v_2' \rightarrowtail_{k_2} \dots \rightarrowtail_{k_{n-1}} v_{n-1}' \rightarrowtail_{k_{n-1}} v_n'$ in G_2 has a pre-image under h in G_1 that starts in v_0: a unique path $\pi : v_0 \rightarrowtail_{k_1} v_1 \rightarrowtail_{k_2} v_2 \rightarrowtail_{k_3} \dots \rightarrowtail_{k_{n-1}} v_{n-1} \rightarrowtail_{k_n} v_n$ in G_1 such that $\pi' = h(\pi)$ holds in the sense of (i).

(iii) $h(V_1) = V_2$, that is, h is surjective.

Proof. Statement (i) can be shown by induction on the length of π, using the (args-forward) property of h from proposition 2.2.11. Analogously, statement (ii) can be established by induction on the length of π', using the (args-backward) property of h from proposition 2.2.11. Statement (iii) follows by applying, for given $v' \in V_2$, the statement of (ii) to an access path π' of v' in G_2 (which exists because term graphs are defined to be root-connected). □

Proposition 2.2.13. Let G, G_1, and G_2 be term graphs over a signature Σ.

(i) If $G \trianglerighteq_h G$, then $h = \mathrm{id}_V$, where id_V is the identity function on the set V of vertices of G.

(ii) If $G_1 \trianglerighteq_h G_2$ and $G_2 \trianglerighteq_g G_1$ hold, then h and g are invertible, $h^{-1} = g$, and consequently $G_1 \sim G_2$.

Proof. For statement (i), suppose that $G \trianglerighteq_h G$ holds for some homomorphism h from G to itself. The fact that $h(v) = \mathrm{id}_V(v)$ holds for all vertices v of G can be established by induction on the length of the shortest access path of v in G. Note that we make use here of the root-connectedness of G (assumed implicitly, see section 2.2) in the form of the assumption that every vertex can be reached by an access path. In order to show statement (ii), note that $G_1 \trianglerighteq_h G_2$ and $G_2 \trianglerighteq_g G_1$ entail $G_1 \trianglerighteq_{h \circ g} G_1$ for homomorphisms h and g from G_1 to G_2. From this $h \circ g = \mathrm{id}_{V_1}$, where V_1 the set of vertices of G_1, follows by (i). Since analogously $h \circ g = \mathrm{id}_{V_2}$ follows, where V_2 is the set of vertices of G_2, the further claims follow. □

§ 2.2.14 (sharing order)**.** The homomorphism relation $\trianglerighteq$ is a preorder on term graphs over a given signature Σ. It induces a partial order on the isomorphism equivalence classes of term graphs over Σ, where anti-symmetry is implied by item proposition 2.2.13 (ii) of the following proposition. We will refer to $\trianglerighteq$ as the *sharing preorder*, and to the induced relation on isomorphism equivalence classes as the *sharing order*.

Notation 2.2.15. Note that, deviating from some literature [51], we use the order relation $\trianglerighteq$ in the same direction as $\leq$: if $G_1 \trianglerighteq G_2$, then G_2 is greater or equal to G_1 with respect to the ordering $\trianglerighteq$ indicating that sharing is typically increased from G_1 to G_2.

Notation 2.2.16 (equivalence classes of term graphs)**.** Let $\mathcal{K} \subseteq \mathrm{TG}(\Sigma)$ be a class of term graphs over signature Σ. For $G \in \mathrm{TG}(\Sigma)$ we will use the notation

$$[\boldsymbol{G}]^{\mathcal{K}}_{\underline{\leftrightarrow}} := \left\{ \boldsymbol{G'} \;\middle|\; G' \in \mathcal{K}, G \leftrightarrows G' \right\}$$

to denote the bisimulation equivalence class of $\boldsymbol{G}$ (the $\sim$-equivalence-class of G) with respect to (the $\sim$-equivalence-classes in) $\mathcal{K}$. And we will write

$$(\boldsymbol{G}\twoheadrightarrow)^{\mathcal{K}} := \left\{ \boldsymbol{G'} \;\middle|\; G' \in \mathcal{K}, G \twoheadrightarrow G' \right\}$$

to denote the class of all $\sim$-equivalence classes in $\mathcal{K}$ that are reachable from $\boldsymbol{G}$ via functional bisimulation. For $\mathcal{K} = \mathrm{TG}(\Sigma)$ we drop the superscript $\mathcal{K}$, and simply write $[\boldsymbol{G}]_{\underline{\leftrightarrow}}$ and $(\boldsymbol{G} \twoheadrightarrow)$.

§ 2.2.17 (complete lattice)**.** A partially ordered set $\langle A, \leq \rangle$ is a *complete lattice* if every subset of A possesses a least upper bound and a greatest lower bound.

Proposition 2.2.18. Let Σ be a signature, and G be a term graph over Σ. The bisimulation equivalence class $[\boldsymbol{G}]_{\underline{\leftrightarrow}}$ of the isomorphism equivalence class $\boldsymbol{G}$ of G is ordered by functional bisimulation $\twoheadrightarrow$ such that $\langle (\boldsymbol{G} \twoheadrightarrow), \twoheadrightarrow \rangle$ is a complete lattice.

Remark 2.2.19. The statement of proposition 2.2.18 is a restriction to sets of $\twoheadrightarrow$-successors of the statements [2, Theorem 3.19] and [51, Theorem 13.2.20], which confer the complete-lattice property for entire bisimulation equivalence classes (of $\sim$-equivalence classes) of term graphs.

Definition 2.2.20 (closedness under functional bisimulation)**.** Let $\mathcal{K} \subseteq \mathrm{TG}(\Sigma)$ be a subclass of the term graphs over some signature Σ. We say that $\mathcal{K}$ is *closed under functional bisimulation* (*closed under bisimulation*), if for all term graphs $G, G' \in \mathrm{TG}(\Sigma)$, whenever $G \in \mathcal{K}$ and $G \twoheadrightarrow G'$ ($G \leftrightarrows G'$), then also $G' \in \mathcal{K}$.

2.3 Introduction

§ 2.3.1 (premise)**.** In this chapter we seek to explore graph formalisms to adequately represent $\lambda_{\mathsf{letrec}}$-terms. In developing these graph representations we draw heavily on ideas from the previous chapter such as λ-transition-graphs as a representation for λ-terms (section 1.9) and the general concepts of including scope$^+$ 'exctended scope' in the formalisms. The inclusion of scope$^+$ is intended to reflect the scoping rules in $\boldsymbol{\lambda}_{\mathsf{letrec}}$, i.e. that variables bound by an abstraction

or by let can only occur underneath their point of definition. Here we consider term graphs built from three kinds of vertices representating applications, abstractions, and variable occurrences, respectively.

§ 2.3.2 (three classes of graph formalisms)**.** In particular we study the following three classes of term graphs:

λ-higher-order-term-graphs (section 2.4) are extensions of first-order term graphs by adding a scope function that assigns a set of vertices, its scope, to every abstraction vertex. There are two variants, one with and one without an edge (a *backlink*) from each variable occurrence to its corresponding abstraction vertex. The class with backlinks is similar to *higher-order term graphs* as defined by Blom in [8], and is in fact an adaptation of that concept to the λ-calculus.

abstraction-prefix based λ-higher-order-term-graphs (section 2.5) abbreviated as λ-ap-ho-term-graphs do not have a scope function but assign, to each vertex v, an abstraction prefix consisting of a word of abstraction vertices that includes those abstractions for which v is in their scope (it actually lists all abstractions for which v is in their *extended scope*, see definition 1.7.9). Abstraction prefixes are aggregated information about the scopes entered so far (much as abstraction prefixes in the decomposition CRSs from the previous chapter; see § 1.2.9, § 1.4.3, definition 1.4.13).

λ-term-graphs with scope delimiters (section 2.7) are plain first-order term graphs intended to represent both classes of higher-order term graphs above, and by extension $\lambda_{\mathsf{letrec}}$-terms. Instead of relying upon added structures for describing scopes, they use scope-delimiter vertices to signify the end of scopes (much as the scope-delimiting steps in the decomposition CRSs from the previous chapter; see § 1.4.5). Variable occurrences as well as scope delimiters may or may not have backlinks to their corresponding abstraction vertices.

§ 2.3.3 (desired properties)**.** We develop these graph formalisms with a number of properties in mind, from which we expect to ensure that the representations are meaningful and useful:

- The term graphs should represent some $\lambda_{\mathsf{letrec}}$-term, and in this sense are not be 'meaningless'.

- Each of these classes induces a notion of functional bisimulation and bisimulation, which preserve the unfolding semantics of the term graphs (and therefore also the unfolding semantics of the $\lambda_{\mathsf{letrec}}$-term they represent).
- Each of these classes induces a sharing order, which reflects the sharing present in the represented $\boldsymbol{\lambda}_{\mathsf{letrec}}$-term.
- We are particularly interested in classes of term graphs that are closed under functional bisimulation, which ensures that when increasing sharing in a term graph it still remains meaningful.

§ 2.3.4 (relating the graph formalisms)**.** It is important to stress that in this chapter we are not going to prove any of properties above that involve in any way $\lambda_{\mathsf{letrec}}$-terms, particularly the former three. Instead we take a leap of faith and trust our intuition (and the authority of [8]) that the λ-higher-order-term-graphs are indeed sound representation of $\lambda_{\mathsf{letrec}}$-terms in that sense. As mentioned before focus on the graph formalisms themselves, and relate them amongst another. Particularly we establish a bijective correspondence between the λ-higher-order-term-graphs and λ-ap-ho-term-graphs, and a correspondence between λ-ap-ho-term-graphs and λ-term-graphs with scope delimiters that is 'almost bijective' (bijective up to the sharing of scope delimiter vertices). We show that all of these correspondences preserve and reflect the sharing order. It is only in chapter 3 that make a connection back to $\lambda_{\mathsf{letrec}}$-terms. In particular, we supply a translation of $\lambda_{\mathsf{letrec}}$-terms to λ-term-graphs with scope delimiters and back and show that the above desired properties do indeed hold. By the correspondences these properties extend also to the higher-order graph formalisms, which validate our conjecture about λ-higher-order-term-graphs.

2.4 λ-higher-order-Term-Graphs

Definition 2.4.1 (signatures Σ^λ, $\Sigma^\lambda_{0_0}$, $\Sigma^\lambda_{0_1}$)**.** By Σ^λ we denote the signature $\{@, \lambda\}$ with $\mathrm{ar}(@) = 2$, and $\mathrm{ar}(\lambda) = 1$. By $\Sigma^\lambda_{0_i}$, for $i \in \{0, 1\}$, we denote the extension $\Sigma^\lambda \cup \{0\}$ of Σ^λ where $\mathrm{ar}(0) = i$. Whether the variable vertices have an outgoing edge depends on the value of i. The intention is to consider two variants of term graphs, one with and one without variable backlinks to their corresponding abstraction vertex.

Definition 2.4.2 (term graphs classes $\mathcal{T}_0$ and $\mathcal{T}_1$)**.** The classes of term graphs over $\Sigma^\lambda_{0_0}$ and $\Sigma^\lambda_{0_1}$ are denoted by $\mathcal{T}_0 := \mathrm{TG}\big(\Sigma^\lambda_{0_0}\big)$ and $\mathcal{T}_1 := \mathrm{TG}\big(\Sigma^\lambda_{0_1}\big)$, respectively.

Notation 2.4.3 (vertex subsets)**.** Let $G = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle$ be a term graph over signature Σ. For $l \in \Sigma$ we denote by $\mathrm{V}(l)$ the set of l-vertices of G, that is, the subset of V consisting of all vertices with label l; more formally, $\mathrm{V}(l) := \{v \in V \mid \mathrm{lab}(v) = l\}$.

A 'λ-higher-order-term-graph' consists of a $\Sigma^\lambda_{0_i}$-term-graph together with a scope function that maps abstraction vertices to their scopes$^+$, which are subsets of the graph's vertices.

Terminology 2.4.4 (scope := scope$^+$)**.** Henceforth when we write 'scope' we mean scope$^+$.

Definition 2.4.5 (scope function for $\Sigma^\lambda_{0_i}$-term-graphs)**.** Let $i \in \{0,1\}$ and $G = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle$ be a $\Sigma^\lambda_{0_i}$-term-graph. A function $\mathrm{Sc} : \mathrm{V}(\lambda) \to \wp(V)$ from λ-vertices of G to vertex sets of G is called a *scope function* for G. Such a function is called *correct* if for all $k \in \{0,1\}$, all vertices $v, w \in V$, and all abstraction vertices $x, y \in \mathrm{V}(\lambda)$ the following holds:

$$\Rightarrow\ r \notin \mathrm{Sc}^-(x) \qquad \text{(root)}$$

$$\Rightarrow\ x \in \mathrm{Sc}(x) \qquad \text{(self)}$$

$$x \in \mathrm{Sc}^-(y) \Rightarrow\ \mathrm{Sc}(x) \subseteq \mathrm{Sc}^-(y) \qquad \text{(nest)}$$

$$v \rightarrowtail_k w\ \wedge\ w \in \mathrm{Sc}^-(x) \Rightarrow\ v \in \mathrm{Sc}(x) \qquad \text{(closed)}$$

$$v \in \mathrm{V}(0) \Rightarrow\ \exists x \in \mathrm{V}(\lambda)\ v \in \mathrm{Sc}^-(x) \qquad (\mathrm{scope}_0)$$

$$v \in \mathrm{V}(0)\ \wedge\ v \rightarrowtail w \Rightarrow \begin{cases} w \in \mathrm{V}(\lambda)\ \wedge \\ v \in \mathrm{Sc}(x) \Leftrightarrow w \in \mathrm{Sc}(x) \end{cases} \qquad (\mathrm{scope}_1)$$

where $\mathrm{Sc}^-(v) := \mathrm{Sc}(v) \setminus \{v\}$. Note that if $i = 0$, then (scope_1) is trivially true and hence superfluous, and if $i = 1$, then (scope_0) is redundant, because it follows from (scope_1).

We say that G *admits* a correct scope function if such a function exists for G.

Definition 2.4.6 (λ-ho-term-graph)**.** Let $i \in \{0,1\}$. A *λ-ho-term-graph* (short for *λ-higher-order-term-graph*) over $\Sigma^\lambda_{0_i}$, is a tuple $\mathcal{G} = \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{Sc}\rangle$ where $G_\mathcal{G} = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle$ is a $\Sigma^\lambda_{0_i}$-term-graph, called the term graph *underlying* $\mathcal{G}$, and Sc is a correct scope function for $G_\mathcal{G}$. The classes of λ-ho-term-graphs over $\Sigma^\lambda_{0_0}$ and $\Sigma^\lambda_{0_1}$ will be denoted by $\mathcal{H}^\lambda_0$ and $\mathcal{H}^\lambda_1$.

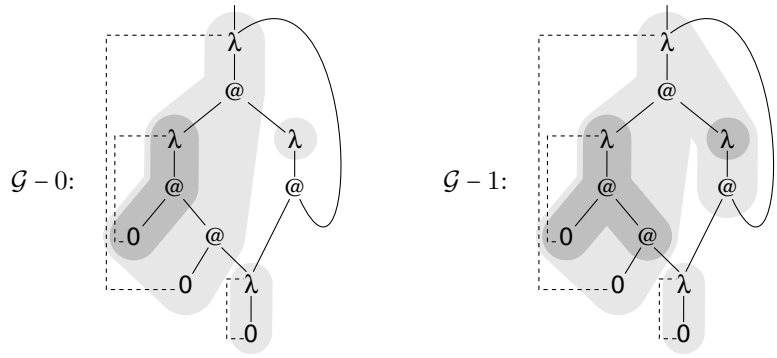

Figure 2.1. $\mathcal{G}_0$ and $\mathcal{G}_1$ are λ-ho-term-graphs in $\mathcal{H}_i^\lambda$ where the dotted backlink edges are present for $i = 1$, but absent for $i = 0$. The underlying term graphs of $\mathcal{G}_0$ and $\mathcal{G}_1$ are identical but their scope functions (signified by the shaded areas) differ. While in $\mathcal{G}_0$ scopes are chosen as small as possible (eager scope-closure), in $\mathcal{G}_1$ some scopes are closed only later in the graph.

Example 2.4.7. Note that there is some freedom on how big the scopes are chosen. The minimal choice corresponds to scope$^+$ in definition 1.7.9. See fig. 2.1 for two different λ-ho-term-graphs over $\Sigma_{0_i}^\lambda$ both of which represent the same term in the λ-calculus with letrec, namely $\mathsf{let}\ f = \lambda x.\,(\lambda y.\,y\ (x\ g))\ (\lambda z.\,g\ f),\ g = \lambda u.\,u\ \mathsf{in}\ f$.

The following lemma states some basic properties of the scope function in λ-ho-term-graphs. Most importantly, scopes in λ-ho-term-graphs are properly nested, in analogy with scopes$^+$ in finite λ-terms.

Lemma 2.4.8. Let $i \in \{0,1\}$, and let $\mathcal{G} = \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{Sc}\rangle$ be a λ-ho-term-graph over $\Sigma_{0_i}^\lambda$. For $v \in V$ and $x \in \mathrm{V}(\lambda)$ we say that x is a *binder for* v if $v \in \mathrm{Sc}(x)$, and we denote by $\mathrm{bds}(v) := \{x \in \mathrm{V}(\lambda) \mid v \in \mathrm{Sc}(x)\}$ the set of binders of v. Then the following statements hold for all $w \in V$ and $v, v_1, v_2 \in \mathrm{V}(\lambda)$:

(i) If $w \in \mathrm{Sc}(v)$, then v is visited on every access path of w, and all vertices on access paths of w after v are in $\mathrm{Sc}^-(v)$. Hence (since $G_\mathcal{G}$ is a term graph, every vertex has an access path) $\mathrm{bds}(w)$ is finite.

(ii) If $\mathrm{Sc}(v_1) \cap \mathrm{Sc}(v_2) \neq \varnothing$ for $v_1 \neq v_2$, then $\mathrm{Sc}(v_1) \subseteq \mathrm{Sc}^-(v_2)$ or $\mathrm{Sc}(v_2) \subseteq \mathrm{Sc}^-(v_1)$. As a consequence, if $\mathrm{Sc}(v_1) \cap \mathrm{Sc}(v_2) \neq \varnothing$, then $\mathrm{Sc}(v_1) \subset \mathrm{Sc}(v_2)$ or $\mathrm{Sc}(v_1) = \mathrm{Sc}(v_2)$ or $\mathrm{Sc}(v_2) \subset \mathrm{Sc}(v_1)$.

(iii) If $\mathrm{bds}(w) \neq \varnothing$, then $\mathrm{bds}(w) = \{v_0, \ldots, v_n\}$ for $v_0, \ldots, v_n \in \mathrm{V}(\lambda)$ and $\mathrm{Sc}(v_n) \subset \mathrm{Sc}(v_{n-1}) \ldots \subset \mathrm{Sc}(v_0)$.

Proof. Let $i \in \{0, 1\}$, and let $\mathcal{G} = \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{Sc}\rangle$ be a λ-ho-term-graph over $\Sigma^{\lambda}_{0_i}$.

For showing (i), let $w \in V$ and $v \in \mathrm{V}(\lambda)$ be such that $w \in \mathrm{Sc}(v)$. Suppose that $\pi : r = w_0 \rightarrowtail_{k_1} w_1 \rightarrowtail_{k_2} w_2 \rightarrowtail_{k_3} \ldots \rightarrowtail_{k_n} w_n = w$ is an access path of w. If $w = v$, then nothing remains to be shown. Otherwise $w_n = w \in \mathrm{Sc}^-(v)$, and, if $n > 0$, then by (closed) it follows that $w_{n-1} \in \mathrm{Sc}(v)$. This argument can be repeated to find subsequently smaller i with $w_i \in \mathrm{Sc}(v)$ and $w_{i+1}, \ldots, w_n \in \mathrm{Sc}^-(v)$. We can proceed as long as $w_i \in \mathrm{Sc}^-(v)$. But since, due to (root), $w_0 = r \notin \mathrm{Sc}^-(v)$, eventually we must encounter an i_0 such that such that $w_{i_0+1}, \ldots, w_n \in \mathrm{Sc}^-(v)$ and $w_{i_0} \in \mathrm{Sc}(v) \setminus \mathrm{Sc}^-(v)$. This implies $w_{i_0} = v$, showing that v is visited on π.

For showing (ii), let $w \in V$ and $v_1, v_2 \in \mathrm{V}(\lambda)$, $v_1 \neq v_2$ be such that $w \in \mathrm{Sc}(v_1) \cap \mathrm{Sc}(v_2)$. Let π be an access path of w. Then it follows by (i) that both v_1 and v_2 are visited on π, and that, depending on whether v_1 or v_2 is visited first on π, either $v_2 \in \mathrm{Sc}^-(v_1)$ or $v_1 \in \mathrm{Sc}^-(v_2)$. Then due to (nest) it follows that either $\mathrm{Sc}(v_2) \subseteq \mathrm{Sc}^-(v_1)$ holds or $\mathrm{Sc}(v_1) \subseteq \mathrm{Sc}^-(v_2)$.

Finally, statement (iii) is an easy consequence of statement (ii). □

Remark 2.4.9 (Comparison to 'higher-order term graphs' [8])**.** The notion of λ-ho-term-graph is an adaptation of the notion of 'higher-order term graph' by Blom [8, Definition 3.2.2] for the purpose of representing finite or infinite λ-terms or cyclic (i.e. strongly regular) λ-terms. In particular, λ-ho-term-graphs over $\Sigma^{\lambda}_{0_1}$ correspond closely to higher-order term graphs over signature Σ^{λ}. But they differ in the following respects:

abstractions: Higher-order term graphs in [8] are graph representations of finite or infinite terms in Combinatory Reduction Systems (CRSs). They typically contain abstraction vertices with label $\square$ that represent CRS abstractions. In contrast, λ-ho-term-graphs have abstraction vertices with label λ that denote λ-abstractions.

signature: Whereas higher-order term graphs in [8] are based on an arbitrary CRS signature, λ-ho-term-graphs only contain the application symbol @ and the variable-occurrence symbol $\mathsf{0}$ in addition to the abstraction symbol λ.

variable backlinks and variable occurrence vertices: In [8] there are no explicit vertices that represent variable occurrences. Instead, variable occurrences

are represented by backlink edges to abstraction vertices. Actually, in the formalisation chosen in [8, Definition 3.2.1], a backlink edge does not directly target the abstraction vertex v it refers to, but ends at a special variant vertex $\bar{v}$ of v. (Every such variant abstraction vertex $\bar{v}$ could be looked upon as a variable vertex that is shared by all edges that represent occurrences of the variable bound by the abstraction vertex v.)

In λ-ho-term-graphs over $\Sigma^{\lambda}_{0_1}$ a variable occurrence is represented by a variable vertex that as outgoing edge has a backlink to the abstraction vertex that binds the occurrence.

conditions on the scope function: While the conditions (root), (self), (nest), and (closed) on the scope function in higher-order term graphs in [8, Definition 3.2.2] correspond directly to the respective conditions in definition 2.4.6, the difference between the condition (scope) there and (scope_1) in definition 2.4.6 reflects the difference described in the previous item.

free variables: Whereas the higher-order term graphs in [8] cater for the presence of free variables, free variables have been excluded from the basic format of λ-ho-term-graphs.

In the following, let $i \in \{0,1\}$ and let $\mathcal{G}_1$ and $\mathcal{G}_2$ be λ-ho-term-graphs over $\Sigma^{\lambda}_{0_i}$ with $\mathcal{G}_k = \langle V_k, \text{lab}_k, \text{args}_k, r_k, \text{Sc}_k \rangle$ for $k \in \{1,2\}$.

Definition 2.4.10 (homomorphism)**.** A *homomorphism* from $\mathcal{G}_1$ to $\mathcal{G}_2$ is a morphism from the structure $\langle V_1, \text{lab}_1, \text{args}_1, r_1, \text{Sc}_1 \rangle$ to the structure $\langle V_2, \text{lab}_2, \text{args}_2, r_2, \text{Sc}_2 \rangle$, i.e. a function $h : V_1 \to V_2$ such that h is a homomorphism from G_1 to G_2, the term graphs underlying $\mathcal{G}_1$ and $\mathcal{G}_2$ respectively, and additionally, for all $v \in \text{V}_1(\lambda)$

$$\bar{h}(\text{Sc}_1(v)) = \text{Sc}_2(h(v)) \tag{2.1}$$

where $\bar{h}$ is the homomorphic extension of h to sets over V_1, i.e. $\bar{h} : \wp(V_1) \to \wp(V_2)$, $A \mapsto \{h(a) \mid a \in A\}$.

If there exists a homomorphism h from $\mathcal{G}_1$ to $\mathcal{G}_2$, then we write $\mathcal{G}_1 \rightrightarrows_h \mathcal{G}_2$ or $\mathcal{G}_2 \leftleftarrows_h \mathcal{G}_1$, or, dropping h as subscript, $\mathcal{G}_1 \rightrightarrows \mathcal{G}_2$ or $\mathcal{G}_2 \leftleftarrows \mathcal{G}_1$.

Definition 2.4.11 (isomorphism)**.** An *isomorphism* from $\mathcal{G}_1$ to $\mathcal{G}_2$ is a homomorphism from $\mathcal{G}_1$ to $\mathcal{G}_2$ that, as a function from V_1 to V_2, is bijective.

If there exists an isomorphism $i : V_1 \to V_2$ from $\mathcal{G}_1$ to $\mathcal{G}_2$, then we say that $\mathcal{G}_1$ and $\mathcal{G}_2$ are *isomorphic*, and write $\mathcal{G}_1 \sim_i \mathcal{G}_2$ or simply $\mathcal{G}_1 \sim \mathcal{G}_2$. The property of

existence of an isomorphism between two λ-ho-term-graphs forms an equivalence relation. We denote by $\boldsymbol{\mathcal{H}}_0^{\boldsymbol{\lambda}}$ and $\boldsymbol{\mathcal{H}}_1^{\boldsymbol{\lambda}}$ the isomorphism equivalence classes of λ-ho-term-graphs over $\Sigma_{0_0}^{\lambda}$ and $\Sigma_{0_1}^{\lambda}$, respectively.

Definition 2.4.12 (bisimulation)**.** A *bisimulation* between $\mathcal{G}_1$ and $\mathcal{G}_2$ is a (λ-ho-term-graph-like) structure $\mathcal{G} = \langle R, \mathrm{lab}, \mathrm{args}, r, \mathrm{Sc}\rangle$ where $\langle R, \mathrm{lab}, \mathrm{args}, r\rangle \in \mathrm{TG}^-(\Sigma)$, $R \subseteq V_1 \times V_2$ and $r = \langle r_1, r_2\rangle$ such that $\mathcal{G}_1 \leftleftarrows_{\pi_1} \mathcal{G} \rightrightarrows_{\pi_2} \mathcal{G}_2$ where π_1 and π_2 are projection functions, defined, for $i \in \{1, 2\}$, by $\pi_i : V_1 \times V_2 \to V_i$, $\langle v_1, v_2\rangle \mapsto v_i$. If there a bisimulation R exists between $\mathcal{G}_1$ and $\mathcal{G}_2$, then we write $\mathcal{G}_1 \leftrightarrows_R \mathcal{G}_2$, or just $\mathcal{G}_1 \leftrightarrows \mathcal{G}_2$.

2.5 Abstraction-prefix based λ-ho-term-graphs

§ 2.5.1 (overview)**.** By an 'abstraction-prefix based λ-higher-order-term-graph' we will mean a term-graph over $\Sigma_{0_i}^{\lambda}$ for $i \in \{0, 1\}$ that is endowed with a correct abstraction prefix function. Such a function P maps abstraction vertices w to words $\mathrm{P}(w)$ consisting of all those abstraction vertices that have w in their scope$^+$, in the order of their nesting from outermost to innermost abstraction vertex. If $\mathrm{P}(w) = v_1 \dots v_n$, then $v_1, \dots, v_n$ are the abstraction vertices that have w in their scope, with v_1 the outermost and v_n the innermost such abstraction vertex.

§ 2.5.2 (comparison to λ-ho-term-graphs)**.** So the conceptual difference between the scope functions of λ-ho-term-graphs defined in the previous section, and abstraction-prefix functions of the λ-ap-ho-term-graphs defined below is the following: A scope function Sc associates with every abstraction vertex v the information on its scope$^+$ $\mathrm{Sc}(v)$ and makes it available at v. In contrast, an abstraction-prefix function P gathers all the scoping information that is relevant to a vertex v (in the sense that it contains all abstraction vertices in whose scope v is) and makes it available at v in the form $\mathrm{P}(v)$. The fact that abstraction-prefix functions make relevant scope information locally available at all vertices leads to simpler correctness conditions.

Definition 2.5.3 (abstraction-prefix function for $\Sigma_{0_i}^{\lambda}$-term-graphs)**.** Let $i \in \{0, 1\}$ and $G = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle$ be a $\Sigma_{0_i}^{\lambda}$-term-graph. A function $\mathrm{P} : V \to V^*$ from vertices of G to words of vertices is called an *abstraction-prefix function*

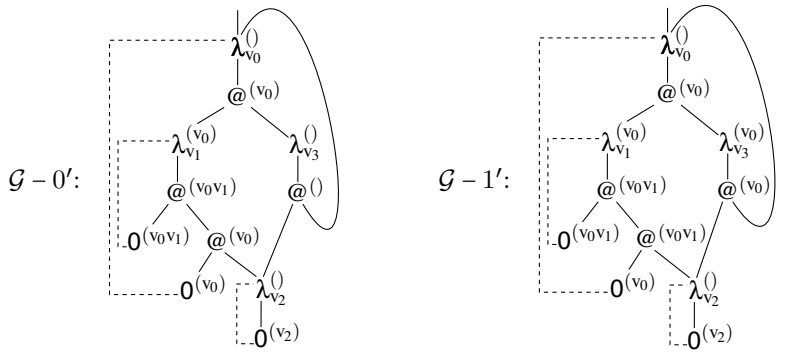

Figure 2.2. The λ-ap-ho-term-graphs corresponding to the λ-ho-term-graphs in fig. 2.1. The subscripts of abstraction vertices indicate their names. The super-scripts of vertices indicate their abstraction-prefixes.

for G. Such a function is called *correct* if for all $v, w \in V$ and $k \in \{0,1\}$:

$$\mathrm{P}(r) = \epsilon \qquad \text{(root)}$$

$$v \in \mathrm{V}(\lambda) \ \wedge \ v \rightarrowtail w \ \Rightarrow \mathrm{P}(w) \leq \mathrm{P}(v)v \qquad (\lambda)$$

$$v \in \mathrm{V}(@) \ \wedge \ v \rightarrowtail w \ \Rightarrow \mathrm{P}(w) \leq \mathrm{P}(v) \qquad (@)$$

$$v \in \mathrm{V}(\mathsf{0}) \ \Rightarrow \mathrm{P}(v) \neq \epsilon \qquad (\mathsf{0}_0)$$

$$v \in \mathrm{V}(\mathsf{0}) \ \wedge \ v \rightarrowtail w \ \Rightarrow w \in \mathrm{V}(\lambda) \ \wedge \ \mathrm{P}(w)\, w = \mathrm{P}(v) \qquad (\mathsf{0}_1)$$

Analogous to definition 2.4.6, if $i = 0$, then $(\mathsf{0}_0)$ is trivially true and hence superfluous, and if $i = 1$, then $(\mathsf{0}_1)$ is redundant, because it follows from $(\mathsf{0}_1)$.

We say that G *admits* a correct abstraction-prefix function if such a function exists for G.

Definition 2.5.4 (λ-ap-ho-term-graph)**.** Let $i \in \{0,1\}$. A *λ-ap-ho-term-graph* (short for *abstraction-prefix based λ-higher-order-term-graph*) over signature $\Sigma^{\lambda}_{\mathsf{0}_i}$ is a tuple $\mathcal{G} = \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{P}\rangle$ where $G_{\mathcal{G}} = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle$ is a $\Sigma^{\lambda}_{\mathsf{0}_i}$-term-graph, called the term graph *underlying* $\mathcal{G}$, and P is a correct abstraction-prefix function for $G_{\mathcal{G}}$. The classes of λ-ap-ho-term-graphs over $\Sigma^{\lambda}_{\mathsf{0}_i}$ will be denoted by $\mathcal{H}^{(\lambda)}_i$.

Example 2.5.5. See fig. 2.2 for two λ-ap-ho-term-graphs, which correspond (see example 2.5.12) to the λ-ho-term-graphs in fig. 2.1.

The following lemma states some basic properties of the scope function in λ-ap-ho-term-graphs.

Lemma 2.5.6. Let $i \in \{0,1\}$ and let $\mathcal{G} = \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{P}\rangle$ be a λ-ap-ho-term-graph over $\Sigma^{\lambda}_{0_i}$. Then the following statements hold:

(i) Suppose that, for some $v, w \in V$, v occurs in $\mathrm{P}(w)$. Then $v \in \mathrm{V}(\lambda)$, occurs in $\mathrm{P}(w)$ only once, and every access path of w passes through v, but does not end there, and thus $w \neq v$. Furthermore it holds that $\mathrm{P}(v)\, v \leq \mathrm{P}(w)$. And conversely, if $\mathrm{P}(w) = p\, v\, q$ for some $p, q \in V^*$, then $\mathrm{P}(v) = p$.

(ii) Vertices in abstraction prefixes are abstraction vertices, and hence P is of the form $\mathrm{P} : V \to (\mathrm{V}(\lambda))^*$.

(iii) For all $v \in \mathrm{V}(\lambda)$ it holds that $v \notin \mathrm{P}(v)$.

(iv) While access paths might end in vertices in $\mathrm{V}(\mathsf{0})$, they pass only through vertices in $\mathrm{V}(\lambda) \cup \mathrm{V}(@)$.

Proof. Let $i \in \{0,1\}$ and let $\mathcal{G} = \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{P}\rangle$ be a λ-ap-ho-term-graph over $\Sigma^{\lambda}_{0_i}$.

For showing (i), let $v, w \in V$ be such that v occurs in $\mathrm{P}(w)$. Suppose further that π is an access path of w. Note that when walking through π the abstraction prefix starts out empty (due to (root)), and is expanded only in steps from vertices $v' \in \mathrm{V}(\lambda)$ (due to (λ), $(@)$, and $(\mathsf{0}_1)$) in which just v' is added to the prefix on the right (due to (λ)). Since v occurs in $\mathrm{P}(w)$, it follows that $v \in \mathrm{V}(\lambda)$, that v must be visited on π, and that π continues after the visit to v. That π is an access path also entails that v is not visited again on π, hence that $w \neq v$ and that v occurs only once in $\mathrm{P}(w)$, and that $\mathrm{P}(v)\, v$, the abstraction prefix of the successor vertex of v on π, is a prefix of the abstraction prefix of every vertex that is visited on π after v.

Statements (ii) and (iii) follow directly from statement (i).

For showing (iv), consider an access path $\pi : r = w_0 \rightarrowtail \ldots \rightarrowtail w_n$ that leads to a vertex $w_n \in \mathrm{V}(\mathsf{0})$. If $i = 0$, then there is no path that extends π properly beyond w_n. So suppose $i = 1$, and let $w_{n+1} \in V$ be such that $w_n \rightarrowtail_0 w_{n+1}$. Then $(\mathsf{0}_1)$ implies that $\mathrm{P}(w_n) = \mathrm{P}(w_{n+1})\, w_{n+1}$, from which it follows by (i) that w_{n+1} is visited already on π. Hence π does not extend to a longer path that is again an access path. □

In the following, let $i \in \{0,1\}$ and let $\mathcal{G}_1$ and $\mathcal{G}_2$ be λ-ap-ho-term-graphs over $\Sigma^{\lambda}_{0_i}$ with $\mathcal{G}_k = \langle V_k, \mathrm{lab}_k, \mathrm{args}_k, r_k, \mathrm{P}_k\rangle$ for $k \in \{1,2\}$.

Definition 2.5.7 (homomorphism)**.** A *homomorphism* from $\mathcal{G}_1$ to $\mathcal{G}_2$ is a morphism from the structure $\langle V_1, \mathrm{lab}_1, \mathrm{args}_1, r_1, \mathrm{P}_1\rangle$ to the structure $\langle V_2, \mathrm{lab}_2, \mathrm{args}_2, r_2, \mathrm{P}_2\rangle$, i.e. a function $h : V_1 \to V_2$ such that h is a homomorphism from G_1 to G_2, the term graphs underlying $\mathcal{G}_1$ and $\mathcal{G}_2$ respectively, and additionally, for all $v \in V_1$

$$h^*(\mathrm{P}_1(v)) = \mathrm{P}_2(h(v))$$

where h^* is the homomorphic extension of h to words over V_1.

We then write $\mathcal{G}_1 \twoheadrightarrow_h \mathcal{G}_2$, or $\mathcal{G}_2 \twoheadleftarrow_h \mathcal{G}_1$. And we write $\mathcal{G}_1 \twoheadrightarrow \mathcal{G}_2$, or for that matter $\mathcal{G}_2 \twoheadleftarrow \mathcal{G}_1$, if there is a homomorphism from $\mathcal{G}_1$ to $\mathcal{G}_2$.

§ 2.5.8 (isomorphism, bisimulation)**.** Analogous to definition 2.4.11 and definition 2.4.12. We denote by $\boldsymbol{\mathcal{H}}_0^{(\boldsymbol{\lambda})}$ and $\boldsymbol{\mathcal{H}}_1^{(\boldsymbol{\lambda})}$ the isomorphism equivalence classes of λ-ho-term-graphs over $\Sigma^{\lambda}_{0_0}$ and $\Sigma^{\lambda}_{0_1}$, respectively.

§ 2.5.9 (relating λ-ho-term-graphs and λ-ap-ho-term-graphs)**.** The following proposition defines mappings between λ-ho-term-graphs and λ-ap-ho-term-graphs by which we establish a bijective correspondence between the two classes. For both directions the underlying λ-term-graph remains unchanged. A_i derives an abstraction-prefix function P from a scope function by assigning to each vertex a word of its binders in the correct nesting order. B_i defines its scope function Sc by assigning to each λ-vertex v the set of vertices that have v in their prefix (along with v since a vertex never has itself in its abstraction prefix).

Proposition 2.5.10 (relating λ-ho-term-graphs and λ-ap-ho-term-graphs)**.** For each $i \in \{0, 1\}$, the mappings A_i and B_i are well-defined between the class of λ-ho-term-graphs over $\Sigma^{\lambda}_{0_i}$ and the class of λ-ap-ho-term-graphs over $\Sigma^{\lambda}_{0_i}$:

$$\begin{aligned}
&\mathrm{A}_i : \mathcal{H}_i^{\lambda} \to \mathcal{H}_i^{(\lambda)}\\
&\quad \mathcal{G} = \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{Sc}\rangle \mapsto \mathrm{A}_i(\mathcal{G}) := \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{P}\rangle\\
&\qquad \text{where } \begin{array}{l}\mathrm{P} : V \to V^*\\ v \mapsto v_0 \dots v_n\end{array} \text{ with } \begin{array}{l}\mathrm{bds}(v) \setminus \{v\} = \{v_0, \dots, v_n\} \text{ and}\\ \mathrm{Sc}(v_n) \subset \mathrm{Sc}(v_{n-1}) \dots \subset \mathrm{Sc}(v_0)\end{array}
\end{aligned}$$

$$\begin{aligned}
&\mathrm{B}_i : \mathcal{H}_i^{(\lambda)} \to \mathcal{H}_i^{\lambda}\\
&\quad \mathcal{G} = \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{P}\rangle \mapsto \mathrm{A}_i(\mathcal{G}) := \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{Sc}\rangle\\
&\qquad \text{where } \begin{array}{l}\mathrm{Sc} : \mathrm{V}(\lambda) \to \wp(V)\\ \qquad v \mapsto \{w \in V \mid v \text{ occurs in } \mathrm{P}(w)\} \cup \{v\}\end{array}
\end{aligned}$$

On the isomorphism equivalence classes $\boldsymbol{\mathcal{H}}_i^{\boldsymbol{\lambda}}$ of $\mathcal{H}_i^{\lambda}$, and $\boldsymbol{\mathcal{H}}_i^{(\boldsymbol{\lambda})}$ of $\mathcal{H}_i^{(\lambda)}$, the functions A_i and B_i induce the functions $\boldsymbol{A}_i : \boldsymbol{\mathcal{H}}_i^{\boldsymbol{\lambda}} \to \boldsymbol{\mathcal{H}}_i^{(\boldsymbol{\lambda})}$, $[\mathcal{G}]_{\sim} \mapsto [\mathrm{A}_i(\mathcal{G})]_{\sim}$ and $\boldsymbol{B}_i : \boldsymbol{\mathcal{H}}_i^{(\boldsymbol{\lambda})} \to \boldsymbol{\mathcal{H}}_i^{\boldsymbol{\lambda}}$, $[\mathcal{G}]_{\sim} \mapsto [\mathrm{B}_i(\mathcal{G})]_{\sim}$.

Theorem 2.5.11 (correspondence between λ-ho-term-graphs and λ-ap-ho-term-graphs)**.** For each $i \in \{0, 1\}$ it holds that the mappings A_i and B_i as defined in proposition 2.5.10 are each other's inverse; thus they define a bijective correspondence between the class of λ-ho-term-graphs over $\Sigma_{0_i}^{\lambda}$ and the class of λ-ap-ho-term-graphs over $\Sigma_{0_i}^{\lambda}$. Furthermore, they preserve and reflect the sharing orders on $\mathcal{H}_i^{\lambda}$ and on $\mathcal{H}_i^{(\lambda)}$:

$$\begin{array}{ll} \forall \mathcal{G}_1, \mathcal{G}_2 \in \mathcal{H}_i^{\lambda} & \mathcal{G}_1 \trianglerighteq \mathcal{G}_2 \iff \mathrm{A}_i(\mathcal{G}_1) \trianglerighteq \mathrm{A}_i(\mathcal{G}_1) \\ \forall \mathcal{G}_1, \mathcal{G}_2 \in \mathcal{H}_i^{(\lambda)} & \mathrm{B}_i(\mathcal{G}_1) \trianglerighteq \mathrm{B}_i(\mathcal{G}_1) \iff \mathcal{G}_1 \trianglerighteq \mathcal{G}_2 \end{array}$$

Example 2.5.12. The λ-ho-term-graphs in fig. 2.1 correspond to the λ-ap-ho-term-graphs in fig. 2.2 via the mappings A_i and B_i as follows:

$$\mathrm{A}_i(\mathcal{G}_0) = \mathcal{G}'_0 \qquad \mathrm{A}_i(\mathcal{G}_1) = \mathcal{G}'_1 \qquad \mathrm{B}_i(\mathcal{G}_0) = \mathcal{G}'_0 \qquad \mathrm{A}_i(\mathcal{G}_1) = \mathcal{G}'_1$$

§ 2.5.13 (higher-order term graphs)**.** Due to this correspondence will henceforth sometimes say 'higher-order term graph' to refer to either λ-ho-term-graphs or λ-ap-ho-term-graphs.

2.6 λ-Term-Graphs without Scope Delimitiers

§ 2.6.1 (Overview)**.** In this section we examine (and dismiss) a naive approach to implement of functional bisimulation on higher-order term graphs as functional bisimulation on their underlying term graph, hoping that this application extends to the higher-order term graph without further ado. We demonstrate that this approach fails, concluding that a faithful first-order implementation of functional bisimulation must not neglect the scoping information that the higher-order term graphs carry.

§ 2.6.2 (variable backlinks)**.** For higher-order term graphs over the signature $\Sigma_{0_0}^{\lambda}$ (i.e. without variable backlinks) essential binding information is lost when looking only at the underlying term graph, to the extent that λ-terms cannot be unambiguously represented anymore. For instance the higher-order term graphs that represent the λ-terms $\lambda xy.\, x\ y$ and $\lambda xy.\, x\ x$ have the same underlying term graph.

This is different for higher-order term graphs over $\Sigma^{\lambda}_{0_1}$, because the abstraction vertex to which a variable vertex belongs is uniquely identified by the variable backlink. This is why in this section we only consider term graphs and higher-order term graphs over $\Sigma^{\lambda}_{0_1}$.

Definition 2.6.3 (λ-term-graphs over $\Sigma^{\lambda}_{0_1}$)**.** A term graph over $\Sigma^{\lambda}_{0_1}$ is called a *λ-term-graph* over $\Sigma^{\lambda}_{0_1}$ if it admits a correct abstraction-prefix function (or equivalently a correct scope function). By $\mathcal{T}^{\lambda}_1$ we denote the class of λ-term-graphs over $\Sigma^{\lambda}_{0_1}$.

Definition 2.6.4 (functional bisimulation on the underlying term graph)**.** Let $\mathcal{G}$ be a higher-order term graph over $\Sigma^{\lambda}_{0_i}$ for $i \in \{0,1\}$ with underlying term graph G. And suppose that there is a homomorphism h from G to G' for some term graph G' over $\Sigma^{\lambda}_{0_i}$.

We say that h *extends to a homomorphism on* $\mathcal{G}$ if there is a higher-order term graph $\mathcal{G}'$ over $\Sigma^{\lambda}_{0_i}$ which has G' as its underlying term graph and h is a homomorphism from $\mathcal{G}$ to $\mathcal{G}'$.

We say that a class $\mathcal{K}$ of higher-order term graphs is *closed under functional bisimulation on the underlying term graphs* if for every $\mathcal{G} \in \mathcal{K}$ with underlying term graph G, every homomorphism h on G extends to a homomorphism on $\mathcal{G}$.

Proposition 2.6.5. Neither the class $\mathcal{H}^{\lambda}_1$ of λ-ho-term-graphs nor the class $\mathcal{H}^{(\lambda)}_1$ of λ-ap-ho-term-graphs is closed under functional bisimulation on the underlying term graphs.

Proof. In view of theorem 2.5.11 it suffices to show the statement for one of the two classes, say $\mathcal{H}^{\lambda}_1$. We show by example that not every homomorphism on the term graph underlying a λ-ho-term-graph over $\Sigma^{\lambda}_{0_1}$ extends to a homomorphism on $\mathcal{H}^{\lambda}_1$.

Consider the λ-ho-term-graphs $\mathcal{G}_1$ and $\mathcal{G}_0$ in fig. 2.3 and their respective underlying term graphs G_1 and G_0. There is an homomorphism h from G_1 to G_0. However, h does not extend to a homomorphism on $\mathcal{G}_1$, since there is no homomorphism from $\mathcal{G}_1$ to $\mathcal{G}_0$, and $\mathcal{G}_0$ is the only λ-ho-term-graph with G_0 as its underlying term graph (G_0 admits only one scope function). □

The next proposition is merely a reformulation of proposition 2.6.5.

Proposition 2.6.6. The scope-forgetful mapping UL, that maps higher-order term graphs to their underlying term graphs, preserves but does not reflect the

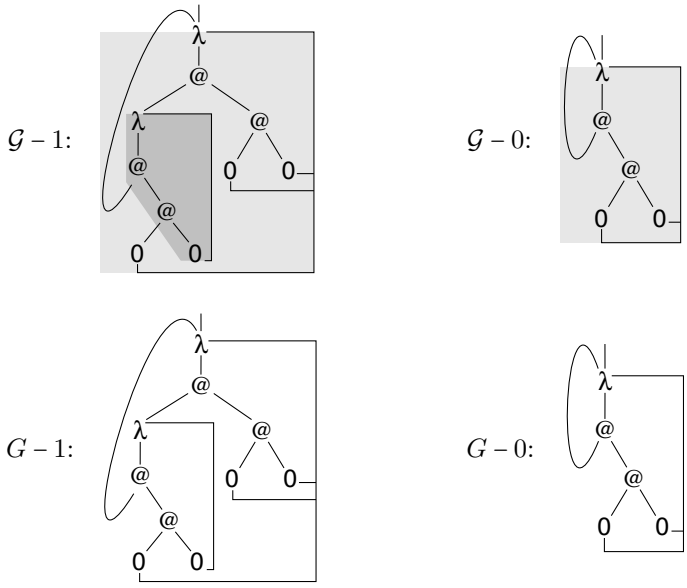

Figure 2.3. There is no homomorphism between the λ-ho-term-graphs $\mathcal{G}_1$ and $\mathcal{G}_0$ because of the scoping condition in definition 2.4.10 (2.1). There is a homomorphism between their underlying term graphs G_1 and G_0 but it does not extend to $\mathcal{G}_1$.

sharing orders on the classes $\mathcal{H}_1^{(\lambda)}$ and $\mathcal{H}_1^{\lambda}$. In particular for $\mathcal{K} \in \{\mathcal{H}_1^{(\lambda)}, \mathcal{H}_1^{\lambda}\}$ and $\mathrm{UL} : \mathcal{K} \to \mathcal{T}_1$ it holds:

$$\begin{aligned} \forall \mathcal{G}_1, \mathcal{G}_2 \in \mathcal{K} \qquad & \mathcal{G}_1 \trianglerighteq \mathcal{G}_2 \;\Rightarrow\; \mathrm{UL}(\mathcal{G}_1) \trianglerighteq \mathrm{UL}(\mathcal{G}_2) \\ \neg\forall \mathcal{G}_1, \mathcal{G}_2 \in \mathcal{K} \qquad & \mathrm{UL}(\mathcal{G}_1) \trianglerighteq \mathrm{UL}(\mathcal{G}_2) \;\Rightarrow\; \mathcal{G}_1 \trianglerighteq \mathcal{G}_2 \end{aligned}$$

§ 2.6.7 (conclusion)**.** As a consequence of this proposition it is not possible to faithfully implement functional bisimulation on higher-order term graphs by only considering their underlying term graphs. It seems that we cannot simply discard the scoping information that is present in the higher-order term graphs. In the next section we refine the above approach by relying on a class of first-order term graphs that accounts for scoping by means of scope delimiter vertices.

2.7 λ-Term-Graphs with Scope Delimiters

§ 2.7.1 (overview)**.** Considering § 2.6.7 we now enrich the signature for λ-term-graphs by an scope-delimiter symbol S (cf. S-steps of the decomposition systems). Delimiter vertices signify the end of a scope. We define variations of that signature where scope-delimiter vertices and/or variable vertices can have backlinks.

§ 2.7.2 (signature for λ-term-graphs with scope delimiters)**.** For all $i \in \{0,1\}$ and $j \in \{1,2\}$ we define the signature $\Sigma^{\lambda}_{0_i,\mathsf{S}_j} := \Sigma^{\lambda} \cup \{0, \mathsf{S}\}$ as an extension of the signature Σ^{λ} where $\mathrm{ar}(0) = i$ and $\mathrm{ar}(\mathsf{S}) = j$, and we denote by $\mathcal{T}_{i,j} := \mathrm{TG}\big(\Sigma^{\lambda}_{0_i,\mathsf{S}_j}\big)$ the class of term graphs over signature $\Sigma^{\lambda}_{0_i,\mathsf{S}_j}$.

§ 2.7.3 (variable and scope-delimiter backlinks)**.** Analogous to the classes $\mathcal{H}^{\lambda}_i$ and $\mathcal{H}^{(\lambda)}_i$, the index i determines whether variable vertices have backlinks to their corresponding abstraction. Here, the additional index i determines whether scope-delimiter vertices have such backlinks (if $j = 2$) or not (if $j = 1$).

§ 2.7.4 (relating scope delimiters and scope)**.** As the scope-delimiter vertices are meant to reproduce the scoping information in the higher-order term graphs, in order to relate the placements of S-vertices and the scoping information we formulate abstraction-prefix conditions for λ-term-graphs (even though they do not carry abstraction prefixes). These conditions are based on definition 2.5.4 and adapted to account for S-vertices which are to decrease the abstraction prefix by one variable.

Definition 2.7.5 (abstraction-prefix function for $\Sigma^{\lambda}_{0_i,\mathsf{S}_j}$-term-graphs)**.** Let $G = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle$ be a $\Sigma^{\lambda}_{0_i,\mathsf{S}_j}$-term-graph for an $i \in \{0,1\}$ and an $j \in \{1,2\}$. A function $\mathrm{P} : V \to V^*$ from vertices of G to words of vertices is called an *abstraction-prefix function* for G. Such a function is called *correct* if for all $v, w \in V$ and $k \in \{0,1\}$ the following holds:

$$\begin{array}{lcll}
 & & \mathrm{P}(r) = \epsilon & (\text{root}) \\
v \in \mathrm{V}(\lambda) \ \wedge \ v \rightarrowtail_0 w & \Rightarrow & \mathrm{P}(w) = \mathrm{P}(v)\, v & (\lambda) \\
v \in \mathrm{V}(@) \ \wedge \ v \rightarrowtail_k w & \Rightarrow & \mathrm{P}(w) = \mathrm{P}(v) & (@) \\
v \in \mathrm{V}(0) & \Rightarrow & \mathrm{P}(v) \neq \epsilon & (0_0) \\
v \in \mathrm{V}(0) \ \wedge \ v \rightarrowtail_0 w & \Rightarrow & w \in \mathrm{V}(\lambda) \ \wedge \ \mathrm{P}(w)\, w = \mathrm{P}(v) & (0_1) \\
v \in \mathrm{V}(\mathsf{S}) \ \wedge \ v \rightarrowtail_0 w & \Rightarrow & \mathrm{P}(w)\, u = \mathrm{P}(v) \text{ for some } u \in V & (\mathsf{S}_0) \\
v \in \mathrm{V}(\mathsf{S}) \ \wedge \ v \rightarrowtail_1 w & \Rightarrow & w \in \mathrm{V}(\lambda) \ \wedge \ \mathrm{P}(w)\, w = \mathrm{P}(v) & (\mathsf{S}_1)
\end{array}$$

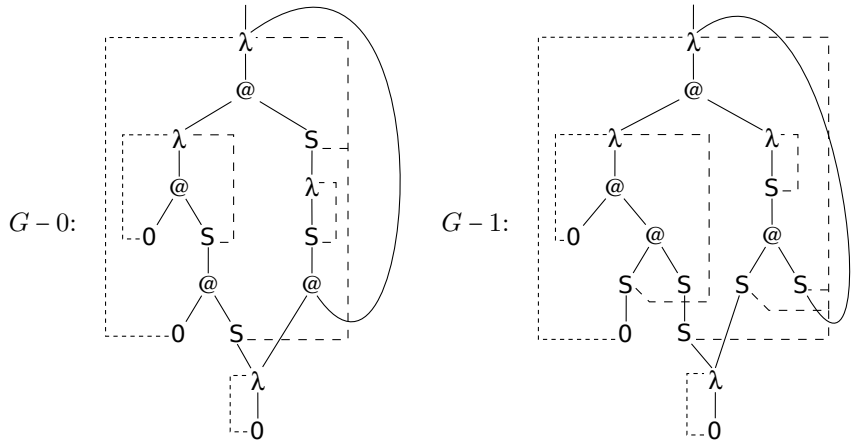

Figure 2.4. The λ-term-graphs corresponding to the λ-ap-ho-term-graphs from fig. 2.2 and the λ-ho-term-graphs from fig. 2.1.

Note that analogous to definition 2.4.6 and definition 2.5.4, if $i = 0$, then $(\mathsf{0}_1)$ is trivially true, and if $i = 1$ then $(\mathsf{0}_0)$ is redundant, because it follows from $(\mathsf{0}_1)$. Additionally, if $j = 1$, then (S_1) is trivially true.

Definition 2.7.6 (λ-term-graphs over $\Sigma^{\lambda}_{\mathsf{0}_i,\mathsf{S}_j}$)**.** Let $i \in \{0,1\}$ and $j \in \{1,2\}$. A *λ-term-graph* over $\Sigma^{\lambda}_{\mathsf{0}_i,\mathsf{S}_j}$ is a $\Sigma^{\lambda}_{\mathsf{0}_i,\mathsf{S}_j}$-term-graph that admits a correct abstraction-prefix function. The class of λ-term-graphs over $\Sigma^{\lambda}_{\mathsf{0}_i,\mathsf{S}_j}$ is denoted by $\mathcal{T}^{\lambda}_{i,j}$, and the class of isomorphism equivalence classes of λ-term-graphs over $\Sigma^{\lambda}_{\mathsf{0}_i,\mathsf{S}_j}$ by $\boldsymbol{\mathcal{T}}^{(\boldsymbol{\lambda})}_{i,j}$.

Example 2.7.7. See fig. 2.4 for examples, that, as we will see, correspond to the ho-term-graphs in fig. 2.1 and in fig. 2.2. A precise formulation of this correspondence is given in example 2.7.21.

The following lemma states some basic properties of λ-term-graphs.

Lemma 2.7.8. Let $i \in \{0,1\}$, and $j \in \{1,2\}$. Let $G = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle$ be a λ-term-graph over $\Sigma^{\lambda}_{\mathsf{0}_i,\mathsf{S}_j}$, and let P be a correct abstraction-prefix function for G. Then the statements (i)–(iii) from lemma 2.5.6 hold as items (i)–(iii) of this lemma, and additionally:

(iv) When stepping through G along a path, P behaves like a stack data structure, in the sense that in a step $v \rightarrowtail w$ it holds:

- if $v \in \mathrm{V}(\lambda)$, then $\mathrm{P}(w) = \mathrm{P}(v)\, v$, that is, $\mathrm{P}(w)$ is obtained by adding v to $\mathrm{P}(v)$ on the right (v is 'pushed on the stack');
- if $v \in \mathrm{V}(\mathsf{0}) \cup \mathrm{V}(\mathsf{S})$, then $\mathrm{P}(w)\, v = \mathrm{P}(v)$ for some $v \in \mathrm{V}(\lambda)$, that is, $\mathrm{P}(w)$ is obtained by removing the rightmost vertex v from $\mathrm{P}(v)$ (v is 'popped from the stack');
- if $v \in \mathrm{V}(@)$, then $\mathrm{P}(w) = \mathrm{P}(v)$, i.e. the abstraction prefix remains the same.

(v) Access paths may end in vertices in $\mathrm{V}(\mathsf{0})$, but only pass through vertices in $\mathrm{V}(\lambda) \cup \mathrm{V}(@) \cup \mathrm{V}(\mathsf{S})$, and depart from vertices in $\mathrm{V}(\mathsf{S})$ only via indexed edges ${}^{\mathsf{S}}\!\!\rightarrowtail_0$.

(vi) There is precisely one correct abstraction-prefix function for G.

Proof. That also here (i)–(iii) from lemma 2.5.6 hold, and that (v) holds, can be shown analogous to the proof of the respective items of lemma 2.5.6. Statement (iv) is easy to check from the definition (see definition 2.7.5 of a correct abstraction-prefix function for a λ-term-graph). For (vi) it suffices to observe that if P is a correct abstraction-prefix function for G, then, for all $v \in V$, the value $\mathrm{P}(v)$ of P at v can be computed by choosing an arbitrary access path π from r to v and using the conditions (λ), (@), and (S_0) to determine in a stepwise manner the values of P at the vertices that are visited on π. Hereby note that in every transition along an edge on π the length of the abstraction prefix only changes by at most 1. □

Terminology 2.7.9. Lemma 2.7.8 (vi) allows us to speak of *the* abstraction-prefix function of a λ-term-graph.

§ 2.7.10 (λ-term-graphs are first-order)**.** Although the requirement of the existence of a correct abstraction-prefix function restricts their possible forms, λ-term-graphs are first-order term graphs, and as such the definitions of homomorphism, isomorphism, and bisimulation for first-order term graphs from section 2.2 apply to them.

§ 2.7.11. The question arises now how a homomorphism h between λ-term-graphs G_1 and G_2 relates their abstraction-prefix functions P_1 and P_2. As it turns out (proposition 2.7.15) P_2 is the 'homomorphic image' of P_1 under h in the sense of the following definition.

Definition 2.7.12 (homomorphic image)**.** Let $G_1 = \langle V_1, \mathrm{lab}_1, \mathrm{args}_1, r_1\rangle$ and $G_2 = \langle V_2, \mathrm{lab}_2, \mathrm{args}_2, r_2\rangle$ be term graphs over $\Sigma^{\lambda}_{0_i,\mathsf{S}_j}$ for some $i \in \{0,1\}$ and $j \in \{1,2\}$, and let P_1 and P_2 be abstraction-prefix functions (not necessarily correct) for G_1 and G_2, respectively. Furthermore, let $h : V_1 \to V_2$ be a homomorphism from G_1 to G_2. We say that P_2 is the *homomorphic image of* P_1 *under* h if it holds:

$$h^* \circ \mathrm{P}_1 = \mathrm{P}_2 \circ h \tag{2.2}$$

where h^* is the homomorphic extension of h to words over V_1.

§ 2.7.13 (homomorphisms between λ-ap-ho-term-graphs)**.** Note that for λ-ap-ho-term-graphs (2.2) (or rather an equivalent thereof) holds by definition (see definition 2.5.7). This is a consequence of the fact that homomorphisms between λ-ap-ho-term-graphs are defined as morphisms between λ-ap-ho-term-graphs when viewed as algebraical structures. As the abstraction-prefix function is a part of λ-ap-ho-term-graphs it has to be respected by morphisms.

§ 2.7.14 (homomorphisms between λ-term-graphs)**.** λ-term-graphs however, do not include an abstraction-prefix function as part of their formalisation. Here the existence of an abstraction-prefix function is merely a mathematical property that distinguishes them from among the term graphs over the same signature. Homomorphisms between λ-term-graphs are defined as morphisms between the structures that underlie their formalisation, i.e. first-order term graphs, and therefore do not by definition respect abstraction-prefix functions.

However, it turns out that, that due the correctness conditions for λ-term-graphs, homomorphisms between λ-term-graphs do in fact respect their abstraction-prefix function:

Proposition 2.7.15. Let G_1 and G_2 be λ-term-graphs, and let P_1 and P_2 be their abstraction-prefix functions, respectively. Suppose that h is a homomorphism from G_1 to G_2. Then P_2 is the homomorphic image of P_1 under h.

This proposition follows from statement (i) of the following lemma. (ii) states that functional bisimulation on term graphs over $\Sigma^{\lambda}_{0_i,\mathsf{S}_j}$ preserves and reflects correctness of the abstraction-prefix functions.

Lemma 2.7.16. Let G_1 and G_2 be term graphs over $\Sigma^{\lambda}_{0_i,\mathsf{S}_j}$ for some $i \in \{0,1\}$ and $j \in \{1,2\}$. Let h be a homomorphism from G_1 to G_2. Furthermore, let P_1 and P_2 be their abstraction-prefix functions, respectively. Then the following two statements hold:

(i) If P_1 and P_2 are correct for G_1 and G_2, respectively, then P_2 is the homomorphic image of P_1.

(ii) If P_2 is the homomorphic image of P_1, then P_1 is correct for G_1 if and only if P_2 is correct for G_2.

Proof. Let $G_1 = \langle V_1, \mathrm{lab}_1, \mathrm{args}_1, r_1\rangle$ and $G_2 = \langle V_2, \mathrm{lab}_2, \mathrm{args}_2, r_2\rangle$. In this proof we denote by $\rightarrowtail'$ the directed-edge relation in G_2, thus if we for instance write $v \rightarrowtail'_0 w$ we mean to say that w is the first child of v in G_2.

For proving (i), we assume that the abstraction-prefix functions P_1 and P_2 are correct for G_1 and G_2, respectively. We establish that P_2 is the homomorphic image under h of P_1 by showing that

$$\forall v \in V_1 \quad h^*(\mathrm{P}_1(v)) = \mathrm{P}_2(h(v)) \tag{2.3}$$

by induction on the length of an access path $\pi : r_1 = v_0 \rightarrowtail v_1 \rightarrowtail \ldots \rightarrowtail v_n = v$ of v in G_1.

If $|\pi| = 0$, then $v = r_1$. It follows from the correctness condition (root) for abstraction-prefix functions and the condition (roots) for homomorphisms that $h^*(\mathrm{P}_1(v)) = h^*(\mathrm{P}_1(r_1)) = h^*(\epsilon) = \epsilon = \mathrm{P}_2(r_2) = \mathrm{P}_2(h(r_1))$.

If $|\pi| = n + 1$, then π is of the form $\pi : r_1 = v_0 \rightarrowtail v_1 \rightarrowtail \ldots \rightarrowtail v_n \rightarrowtail_i v_{n+1} = v$ for some $i \in \{0, 1\}$. We have to show that $h^*(\mathrm{P}_1(v)) = \mathrm{P}_2(h(v))$ holds, with $h^*(\mathrm{P}_1(v_n)) = \mathrm{P}_2(h(v_n))$ as an induction hypothesis. We will do so by distinguishing the three possible labels v_n can have, namely λ, @, and S (see lemma 2.7.8 (v)) and by applying the correctness conditions from definition 2.7.5.

$v_n \in \mathrm{V}_1(\lambda)$: Since h is a homomorphism also $h(v_n) \in \mathrm{V}_2(\lambda)$ holds. Applying the correctness condition (λ) to v_n and $h(v_n)$ yields that $\mathrm{P}_1(v) = \mathrm{P}_1(v_n)\, v_n$ and $\mathrm{P}_2(h(v)) = \mathrm{P}_2(h(v_n))\, h(v_n)$. From this we now obtain $h^*(\mathrm{P}_1(v)) = h^*(\mathrm{P}_1(v_n)\, v_n) = h^*(\mathrm{P}_1(v_n))\, h(v_n) = \mathrm{P}_2(h^*(v_n))\, h(v_n) = \mathrm{P}_2(h(v))$ by using the induction hypothesis.

$v_n \in \mathrm{V}_1(@)$: Since h is a homomorphism also $h(v_n) \in \mathrm{V}_2(@)$ holds. Applying the correctness condition (@) to v_n and $h(v_n)$ yields that $\mathrm{P}_1(v) = \mathrm{P}_1(v_n)$ and $\mathrm{P}_2(h(v)) = \mathrm{P}_2(h(v_n))$. Then we obtain $h^*(\mathrm{P}_1(v)) = h^*(\mathrm{P}_1(v_n)) = \mathrm{P}_2(h(v_n)) = \mathrm{P}_2(h(v))$ by using the induction hypothesis.

$v_n \in \mathrm{V}_1(\mathsf{S})$: Since h is a homomorphism also $h(v_n) \in \mathrm{V}_2(\mathsf{S})$ holds. We distinguish two cases for the last step of π (is it via the backlink or not?):

$v_n \rightarrowtail_1 v$: This implies $h(v_n) \rightarrowtail'_1 h(v)$. Applying the correctness condition (S_1) to v_n and to $h(v_n)$ yields $\mathrm{P}_1(v_n) = \mathrm{P}_1(v)\,v$ and $\mathrm{P}_2(h(v))\,h(v) = \mathrm{P}_2(h(v_n))$. Applying the induction hypothesis we get $\mathrm{P}_2(h(v))\,h(v) = \mathrm{P}_2(h(v_n)) = h^*(\mathrm{P}_1(v_n)) = h^*(\mathrm{P}_1(v)\,v) = h^*(\mathrm{P}_1(v))\,h(v)$. From this we conclude $\mathrm{P}_2(h(v)) = h^*(\mathrm{P}_1(v))$.

$v_n \rightarrowtail_0 v$: Analogously, but we rely on (S_0) instead of (S_1).

For showing statement (ii), we assume that – as we have just shown – P_2 is the homomorphic image of P_1.

For the direction "⇒" of the equivalence in (i) we assume that P_1 is correct for G_1, and we show that P_2 is correct for G_2, according to the conditions (root), (λ), (@), $(\mathsf{0}_0)$, and $(\mathsf{0}_1)$ from definition 2.7.5.

For the two conditions that do not involve transitions this is easy to show: The condition (root) for G_2 follows from the condition (roots) from definition 2.2.7 and (2.2). And the condition $(\mathsf{0}_0)$ follows similarly by using that every pre-image under h of a variable vertex in G' must be a variable vertex in G since h is a homomorphism.

Let us only look at one of the remaining cases, and look at the condition (S_0). For this, let $v', v'_0 \in V_2$ such that $v' \in \mathrm{V}_2(\mathsf{S})$ with $v' \rightarrowtail'_0 v'_0$. Then as a consequence of the (args-forward) condition for h from proposition 2.2.11 there exist $v, v_0 \in V_1$ with $v \in \mathrm{V}_1(\mathsf{S})$, $v \rightarrowtail_0 v_0$, and with $h(v) = v'$, $h(v_0) = v'_0$. Since (S_0) is satisfied for G_1, there exists a vertex $w \in V_1$ such that $\mathrm{P}_1(v_0)w = \mathrm{P}_1(v)$. From this and by (2.2) we obtain $\mathrm{P}_2(v'_0)\,h(w) = \mathrm{P}_2(h(v_0))\,h(w) = h^*(\mathrm{P}_1(v_0))\,h(w) = h^*(\mathrm{P}_1(v_0)\,w) = h^*(\mathrm{P}_1(v)) = \mathrm{P}_2(h(v)) = \mathrm{P}_2(v')$. This shows the existence of $w' \in V_2$ (to wit $w' := h(w)$) with $\mathrm{P}_2(v'_0)\,w' = \mathrm{P}_2(v')$. In this way we have established the correctness condition (S_0) for G_2.

The direction "⇐" of the equivalence in (i) can be established analogously by recognizing that the correctness conditions from definition 2.7.5 carry over also from G_2 to G_1 via h due to the homomorphic image property (2.2). The arguments for the individual correctness conditions are analogous to the ones used above, but they depend on using the (args-backward) property from proposition 2.2.11 of h. □

§ 2.7.17 (relating λ-term-graphs defined above and λ-ap-ho-term-graphs)**.** We now proceed to define a precise relationship between λ-term-graphs and λ-ap-ho-term-graphs via mappings that translate between these classes:

The mapping $\mathrm{G}_{i,j}$ produces a λ-term-graph for any given λ-ap-ho-term-graph by adding to the original set of vertices a number of delimiter vertices at

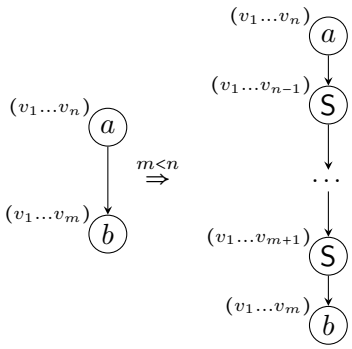

Figure 2.5. Definition of $\mathrm{G}_{i,j}$ by inserting S-vertices, between edge-connected vertices of a λ-ap-ho-term-graph

the appropriate places. That is, at every position where the abstraction prefix decreases by n elements, n S-vertices are inserted as depicted in fig. 2.5. In the image, the original abstraction prefix is retained as part of the vertices, which we use for defining the edges of the image.

The mapping $\mathcal{G}_{i,j}$ back to λ-ap-ho-term-graphs is simpler because it only has to erase the S-vertices, and add the correct abstraction prefix that we already know to exists for the λ-term-graph.

Proposition 2.7.18 (from λ-ap-ho-term-graphs to λ-term-graphs)**.** Let $i \in \{0,1\}$ and $j \in \{1,2\}$. The mapping $\mathrm{G}_{i,j}$ defined below is well-defined between the class of λ-term-graphs over $\Sigma^{\lambda}_{0_i,\mathsf{S}_j}$ and the class of λ-ap-ho-term-graphs over $\Sigma^{\lambda}_{0_i}$:

$$\mathrm{G}_{i,j} : \mathcal{H}^{(\lambda)}_i \to \mathcal{T}^{\lambda}_{i,j},\ \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{P}\rangle \mapsto \left\langle V', \mathrm{lab}', \mathrm{args}', r'\right\rangle$$

where:

$$V' := \{\langle v, \mathrm{P}(v)\rangle \mid v \in V\} \cup \left\{ \langle v, k, v', p\rangle \;\middle|\; v, v' \in V, v \rightarrowtail_k v',\ \begin{array}{l} \mathrm{V}(v) = \lambda \wedge \mathrm{P}(v') < p \leq \mathrm{P}(v)\, v \\ \vee\ \mathrm{V}(v) = @ \wedge \mathrm{P}(v') < p \leq \mathrm{P}(v) \end{array} \right\}$$

$$r' := \langle r, \epsilon\rangle \qquad \mathrm{lab}' : V' \to \Sigma^{\lambda}_{0_i,\mathsf{S}_j},\ \begin{array}{l} \langle v, \mathrm{P}(v)\rangle \mapsto \mathrm{lab}(v) \\ \langle v, k, v', p\rangle \mapsto \mathsf{S} \end{array}$$

and $\mathrm{args}' : V' \to (V')^*$ is defined such that for the induced indexed successor relation $\rightarrowtail'_{(\cdot)}$ it holds:

$$
\begin{aligned}
&v \rightarrowtail_k w \;\wedge\; \#\mathrm{del}(v,k) = 0 \;\Rightarrow\; \langle v, \mathrm{P}(v)\rangle \rightarrowtail'_k \langle w, \mathrm{P}(w)\rangle \\
&v \rightarrowtail_0 w \;\wedge\; \#\mathrm{del}(v,0) > 0 \;\wedge\; \mathrm{lab}(v) = \lambda \;\wedge\; \mathrm{P}(v) = \mathrm{P}(w)\,u\,p \\
&\quad\Rightarrow\; \langle v, \mathrm{P}(v)\rangle \rightarrowtail'_0 \langle v, 0, w, \mathrm{P}(v)\,v\rangle \;\wedge\; \langle v, 0, w, \mathrm{P}(w)\,u\rangle \rightarrowtail'_0 \langle w, \mathrm{P}(w)\rangle \\
&v \rightarrowtail_k w \;\wedge\; \#\mathrm{del}(v,k) > 0 \;\wedge\; \mathrm{lab}(v) = @ \;\wedge\; \mathrm{P}(v) = \mathrm{P}(w)\,u\,p \\
&\quad\Rightarrow\; \langle v, \mathrm{P}(v)\rangle \rightarrowtail'_k \langle v, k, w, \mathrm{P}(v)\rangle \;\wedge\; \langle v, k, w, \mathrm{P}(w)\,u\rangle \rightarrowtail'_0 \langle w, \mathrm{P}(w)\rangle \\
&v \rightarrowtail_k w \;\wedge\; \#\mathrm{del}(v,k) > 0 \;\wedge\; \langle v, k, w, p\,u\rangle, \langle v, k, w, p\rangle \in V' \\
&\quad\Rightarrow\; \langle v, k, w, p\,u\rangle \rightarrowtail'_0 \langle v, k, w, p\rangle \\
&v \rightarrowtail_k w \;\wedge\; \#\mathrm{del}(v,k) > 0 \;\wedge\; \langle v, k, w, p\,u\rangle \in V' \;\wedge\; j = 2 \\
&\quad\Rightarrow\; \langle v, k, w, p\,u\rangle \rightarrowtail'_1 \langle w, \mathrm{P}(w)\rangle
\end{aligned}
$$

for all $v, w, u \in V$, $k \in \{0,1\}$, $p \in V^*$, and where the function #del is defined as:

$$
\#\mathrm{del}(v,k) := \begin{cases} |\mathrm{P}(v)| - |\mathrm{P}(v')| & \text{if } v \in \mathrm{V}(@) \wedge v \rightarrowtail_k v' \\ |\mathrm{P}(v)| + 1 - |\mathrm{P}(v')| & \text{if } v \in \mathrm{V}(\lambda) \wedge v \rightarrowtail_k v' \\ 0 & \text{otherwise} \end{cases}
$$

$\mathrm{G}_{i,j}$ induces the function $\mathbf{G}_{i,j} : \boldsymbol{\mathcal{H}}_i^{(\boldsymbol{\lambda})} \to \boldsymbol{\mathcal{T}}_{i,j}^{(\boldsymbol{\lambda})}$, $\boldsymbol{\mathcal{G}} = [\mathcal{G}]_\sim \mapsto [\mathrm{G}_{i,j}(\mathcal{G})]_\sim$ on the respective isomorphism equivalence classes.

Proposition 2.7.19 (from λ-term-graphs to λ-ap-ho-term-graphs)**.** Let $i \in \{0,1\}$ and $j \in \{1,2\}$. The mapping $\mathcal{G}_{i,j}$ defined below is well-defined between the class of λ-term-graphs over $\Sigma^{\lambda}_{0_i,\mathsf{S}_j}$ and the class of λ-ap-ho-term-graphs over $\Sigma^{\lambda}_{0_i}$:

$$
\mathcal{G}_{i,j} : \mathcal{T}^{\lambda}_{i,j} \to \mathcal{H}^{(\lambda)}_i,\; G = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle \mapsto \langle V', \mathrm{lab}|_{V'}, \mathrm{args}', r, \mathrm{P}'\rangle
$$

where

$$
V' := \mathrm{V}(\lambda) \cup \mathrm{V}(@) \cup \mathrm{V}(\mathsf{0})
$$

$$
\mathrm{args}' : V' \to (V')^*, \text{ so that for the induced indexed successor relation } \rightarrowtail'_{(\cdot)} \text{ on } V' \text{ it holds:}
$$

$$
v \rightarrowtail'_k w \;\Leftrightarrow\; v \rightarrowtail_k \cdot \left({}^{\mathsf{S}}\!\rightarrowtail_0\right)^* w \quad (\forall v, w \in V', k \in \{0,1\})
$$

$$
\mathrm{P}' := \mathrm{P}|_{V'} \text{ where P is the abstraction-prefix function of } G.
$$

$\mathcal{G}_{i,j}$ induces the function $\boldsymbol{\mathcal{G}}_{i,j} : \boldsymbol{\mathcal{T}}^{(\boldsymbol{\lambda})}_{i,j} \to \boldsymbol{\mathcal{H}}^{(\boldsymbol{\lambda})}_i$, $\boldsymbol{G} = [G]_\sim \mapsto [\mathcal{G}_{i,j}(G)]_\sim$ on the isomorphism equivalence classes.

The mappings $\mathcal{G}_{i,j}$ and $\mathrm{G}_{i,j}$ now define a correspondence between the class $\mathcal{H}_i^{(\lambda)}$ of λ-ap-ho-term-graphs and the class $\mathcal{T}_{i,j}^{\lambda}$ of λ-term-graphs in then sense as stated by the following theorem.

Theorem 2.7.20 (correspondence between λ-ap-ho-term-graphs and λ-term-graphs)**.** Let $i \in \{0,1\}$ and $j \in \{1,2\}$. The mappings $\mathcal{G}_{i,j}$ from proposition 2.7.19 and $\mathrm{G}_{i,j}$ from proposition 2.7.18 define a correspondence between the class of λ-term-graphs over $\Sigma^{\lambda}_{0_i,\mathsf{S}_j}$ and the class of λ-ap-ho-term-graphs over $\Sigma^{\lambda}_{0_i}$ with the following properties:

(i) $\mathcal{G}_{i,j} \circ \mathrm{G}_{i,j} = \mathrm{id}_{\mathcal{H}_i^{(\lambda)}}$.

(ii) For all $G \in \mathcal{T}_{i,j}^{\lambda}$: $(\mathrm{G}_{i,j} \circ \mathcal{G}_{i,j})(G) \rightrightarrows^{\mathsf{S}} G$.

(iii) $\mathcal{G}_{i,j}$ and $\mathrm{G}_{i,j}$ preserve and reflect the sharing orders on $\mathcal{H}_i^{(\lambda)}$ and on $\mathcal{T}_{i,j}^{\lambda}$:

$$\begin{aligned} &\forall \mathcal{G}_1, \mathcal{G}_2 \in \mathcal{H}_i^{(\lambda)} && \mathrm{G}_{i,j}(\mathcal{G}_1) \trianglerighteq \mathrm{G}_{i,j}(\mathcal{G}_2) \iff \mathcal{G}_1 \trianglerighteq \mathcal{G}_2 \\ &\forall G_1, G_2 \in \mathcal{T}_{i,j}^{\lambda}) && \mathcal{G}_{i,j}(G_1) \trianglerighteq \mathcal{G}_{i,j}(G_2) \iff G_1 \trianglerighteq G_2 \end{aligned}$$

Furthermore, statements analogous to (i) (ii), and (iii) hold for the correspondences $\mathbf{G}_{i,j}$ and $\boldsymbol{\mathcal{G}}_{i,j}$, induced by $\mathrm{G}_{i,j}$ and $\mathcal{G}_{i,j}$, between the classes $\boldsymbol{\mathcal{H}}_i^{(\boldsymbol{\lambda})}$ and $\boldsymbol{\mathcal{T}}_{i,j}^{(\boldsymbol{\lambda})}$ of isomorphism equivalence classes of graphs in $\boldsymbol{\mathcal{H}}_i^{(\boldsymbol{\lambda})}$ and $\boldsymbol{\mathcal{T}}_{i,j}^{(\boldsymbol{\lambda})}$, respectively.

Example 2.7.21. The λ-ap-ho-term-graphs in fig. 2.2 correspond to the λ-ap-ho-term-graphs in fig. 2.4 via $\mathrm{G}_{i,j}$ and $\mathcal{G}_{i,j}$ as follows:

$$\mathrm{G}_{i,j}(\mathcal{G}_0) = G'_0 \qquad \mathrm{G}_{i,j}(\mathcal{G}_1) = G'_1 \qquad \mathcal{G}_{i,j}(G_0) = \mathcal{G}'_0 \qquad \mathcal{G}_{i,j}(G_1) = \mathcal{G}'_1$$

Proposition 2.7.22. The mapping $\mathcal{G}_{i,j}$ from λ-term-graphs to λ-ap-ho-term-graphs is not injective.

Proof. This is witnessed by the following example. □

Example 2.7.23 ($\mathcal{G}_{i,j}$ is not injective)**.** Consider the λ-ap-ho-term-graph $\mathcal{G} \in \mathcal{H}_0^{(\lambda)}$, and the λ-term-graphs $G, G' \in \mathcal{T}_{0,1}^{\lambda}$ below. It holds that $\mathcal{G}_{0,1}(G) = \mathcal{G} = \mathcal{G}_{0,1}(G')$.

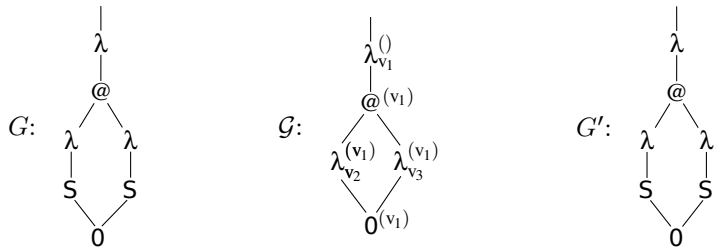

§ 2.7.24 (a weaker than bijective correspondence)**.** In contrast to the correspondence between λ-ho-term-graphs and λ-ap-ho-term-graphs, which is bijective (theorem 2.5.11), due to proposition 2.7.22 we have no bijective correspondence between λ-ap-ho-term-graph and λ-term-graphs. However, note that in example 2.7.23 it holds that $G' \rightrightarrows^{\mathsf{S}} G$, and consequently $G \leftrightarrows^{\mathsf{S}} G'$. We can say that G and G' only differ in their 'degree of S-sharing'. Since $\mathcal{G}_{i,j}$ ignores S-vertices and their sharing, the degree of S-sharing cannot be reflected in the corresponding λ-ap-ho-term-graph. This observation leads us to look for a weaker (than bijective) correspondence between λ-ap-ho-term-graph and λ-term-graphs: the weakening consists of equating S-bisimilar λ-term-graphs as in the following proposition.

Proposition 2.7.25. Let $i \in \{0,1\}$ and $j \in \{1,2\}$. The mapping $\mathcal{G}_{i,j}$ in proposition 2.7.19 maps two λ-term-graphs that are S-bisimilar to the same λ-ap-ho-term-graph. That is, for all λ-term-graphs G_1, G_2 over $\Sigma^{\lambda}_{0_i,\mathsf{S}_j}$ it holds:

$$G_1 \leftrightarrows^{\mathsf{S}} G_2 \implies \mathcal{G}_{i,j}(G_1) \sim \mathcal{G}_{i,j}(G_2)$$

Remark 2.7.26 (λ-term-graphs over $\Sigma^{\lambda}_{0_i,\mathsf{S}_j}$ up to S-bisimilarity)**.** In the original paper [24, 5.15], we use this insight to develop a variation on $\mathcal{T}^{\lambda}_{i,j}$, where the correctness condition from definition 2.7.5 is relaxed to non-delimiter vertices. Consequently an S-vertex is allowed to be a delimiter for two different scopes (as G in example 2.7.23). We show that these λ-term-graphs are always S-bisimilar to graphs in $\mathcal{T}^{\lambda}_{i,j}$ and that there is a bijective correspondence with λ-ap-ho-term-graphs if we consider equivalence classes of these λ-term-graphs up to isomorphism and S-bisimilarity. We omit this part this thesis. Instead we lean on the not quite bijective correspondence from proposition 2.7.25 as is.

§ 2.7.27 (outlook)**.** Be reminded that we seek to implement bisimulation and functional bisimulation on higher-order term graphs via (functional) bisimulation on λ-term-graphs (which are first-order term graphs). Now that we have a

correspondence result that relates λ-term-graphs and higher-order term graphs, we investigate how λ-term-graphs behave under bisimulation and functional bisimulation, specifically which classes of λ-term-graphs are closed under (functional) bisimulation.

2.8 Not closed under (functional) bisimulation

§ 2.8.1 (overview)**.** In this section we collect negative results concerning closedness under bisimulation and functional bisimulation for the classes of λ-term-graphs as introduced in the previous section.

Proposition 2.8.2. None of the classes $\mathcal{T}^{\lambda}_{1}$ and $\mathcal{T}^{\lambda}_{i,j}$, for $i \in \{0,1\}$ and $j \in \{1,2\}$, of λ-term-graphs are closed under bisimulation.

§ 2.8.3. This proposition is an immediate consequence of the following proposition, which can be viewed as a refinement, because it formulates non-closedness of classes of λ-term-graphs under specialisations of bisimulation, namely for functional bisimulation (under which some classes are not closed), and for converse functional bisimulation (under which none of the classes considered here is closed).

Proposition 2.8.4. None of the classes $\mathcal{T}^{\lambda}_{i,j}$ of λ-term-graphs for $i \in \{0,1\}$ and $j \in \{1,2\}$ are closed under functional bisimulation, or under converse functional bisimulation. Additionally, the class $\mathcal{T}^{\lambda}_{1}$ of λ-term-graphs is not closed under converse functional bisimulation.

Proof. We proof the following statements by giving counterexamples:

(i) None of the classes $\mathcal{T}^{\lambda}_{0,j}$ for $j \in \{1,2\}$ are closed under $\twoheadrightarrow$, or under $\twoheadleftarrow$.

(ii) None of the classes $\mathcal{T}^{\lambda}_{1}$ and $\mathcal{T}^{\lambda}_{1,j}$ for $j \in \{1,2\}$ of λ-term-graphs are closed under converse functional bisimulation $\twoheadleftarrow$.

(iii) The class $\mathcal{T}^{\lambda}_{1,1}$ of λ-term-graphs is not closed $\twoheadrightarrow$.

(iv) The class $\mathcal{T}^{\lambda}_{1,2}$ of λ-term-graphs is not closed under $\twoheadrightarrow$.

The counterexamples:

(i) Let Δ be one of the signatures $\Sigma^{\lambda}_{0_0,\mathsf{S}_j}$ $j \in \{1,2\}$. Consider the following term graphs over Δ:

$G-2$: 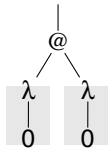

$G-1$: 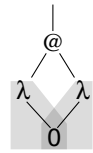

$G-0$:

It holds that $G_2 \trianglerighteq G_1 \trianglerighteq G_0$. But while G_2 and G_0 admit correct abstraction-prefix functions over Δ (since the implied scopes (drawn shaded) are nested), this is not the case for G_1 (overlapping scopes). Therefore, the class of λ-term-graphs over Δ is closed neither under functional bisimulation nor under converse functional bisimulation.

(ii) Let Δ be one of the signatures $\Sigma^{\lambda}_{0_1}$ and $\Sigma^{\lambda}_{0_1,\mathsf{S}_j}$. Consider the following term graphs over Δ:

$G-1$: 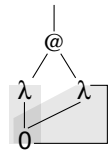

$G-0$: 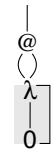

It holds that $G_1 \trianglerighteq G_0$. But while G_0 admits a correct abstraction-prefix function, this is not the case for G_1 (overlapping scopes). Therefore, the class of λ-term-graphs over Δ is not closed under converse functional bisimulation.

(iii) Consider the following term graphs over $\Sigma^{\lambda}_{0_1,\mathsf{S}_1}$:

$G-1$: 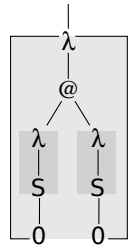

$G-0$: 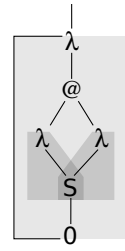

It holds that $G_1 \trianglerighteq G_0$. However, while G_1 admits a correct abstraction-prefix function, this is not the case for G_0 (overlapping scopes). Therefore the class of λ-term-graphs over $\Sigma^{\lambda}_{0_1,\mathsf{S}_1}$ is not closed under functional bisimulation.

(iv) Consider the following term graphs over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$:

$G-1$: 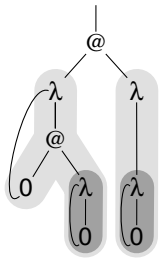

$G-0$: 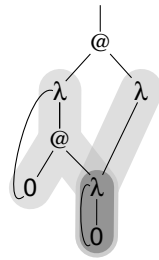

It holds that $G_1 \rightrightarrows G_0$. However, while G_1 admits a correct abstraction-prefix function, this is not the case for G_0 (overlapping scopes). Therefore the class of λ-term-graphs over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ is not closed under functional bisimulation.

□

As an easy consequence of proposition 2.8.2, and of proposition 2.8.4 (i) and (ii), together with the examples used in the proof, we obtain the following two propositions.

Proposition 2.8.5. Let $i \in \{0,1\}$. None of the classes $\mathcal{H}^{\lambda}_i$ of λ-ho-term-graphs, or $\mathcal{H}^{(\lambda)}_i$ of λ-ap-ho-term-graphs are closed under bisimulation on the underlying term graphs.

Proposition 2.8.6. The following statements hold:

(i) Neither $\mathcal{H}^{\lambda}_0$ nor $\mathcal{H}^{(\lambda)}_0$ is closed under functional bisimulation or converse functional bisimulation on the underlying term graphs.

(ii) Neither $\mathcal{H}^{\lambda}_1$ nor $\mathcal{H}^{(\lambda)}_1$ is closed under converse functional bisimulation on the underlying term graphs.

Remark 2.8.7. Note that proposition 2.8.6, (i) is a strengthening of the statement of proposition 2.6.5 earlier.

2.9 Closed under functional bisimulation

§ 2.9.1 (recapitulation)**.** The negative results gathered in the last section leave the impression that our enterprise in a quite poor state: For the classes of λ-term-graphs we introduced, proposition 2.8.4 leaves $\mathcal{T}^{\lambda}_1$ as the only class

of λ-term-graphs which still might be closed under functional bisimulation. Actually, $\mathcal{T}_1^\lambda$ *is* closed (we do not prove this here), but that does not help us any further, because the correspondences in theorem 2.7.20 do not apply to this class, and worse still, proposition 2.6.6 rules out simple correspondences for $\mathcal{T}_1^\lambda$. So in this case we are left without any satisfying correspondences to higher-order term graphs that we have for the other classes of λ-term-graphs; those however are not closed under functional bisimulation.

§ 2.9.2 (eager scope-closure)**.** But in this section we establish that the class $\mathcal{T}_{1,2}^\lambda$ is very useful after all: we find that its restriction to λ-term-graphs with eager scope-closure (definition 2.9.4 below) is in fact closed under functional bisimulation.

Example 2.9.3 (eager-scope λ-term-graphs)**.** Let us look at two $\mathcal{T}_{1,2}^\lambda$-term-graphs from earlier and see whether we can fix closedness under functional bisimulation in that instance by making them eager-scope. Consider the λ-term-graph G_1 from the proof of proposition 2.8.4 (iv). The scopes of the two topmost abstractions are not closed on the paths to variable occurrences belonging to the bottommost abstractions, although both scopes could have been closed immediately before the bottommost abstractions. When this is actually done, and the following variation $\tilde{G}_1$ of G_1 with eager scope-closure is obtained, then the problem disappears:

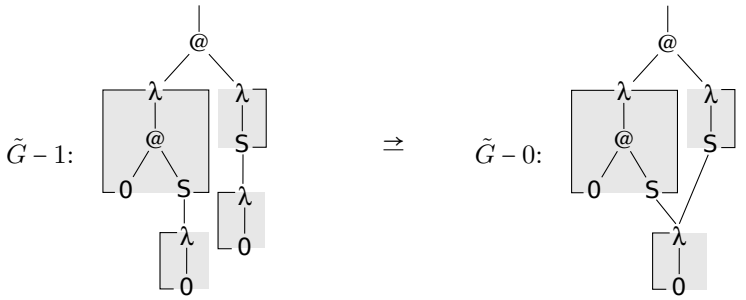

$\tilde{G}_0$ has again a correct abstraction-prefix function and is therefore a λ-term-graph.

For λ-term-graphs over $\Sigma^\lambda_{0_1,\mathsf{S}_1}$ and $\Sigma^\lambda_{0_1,\mathsf{S}_2}$ we define the eager-scope property.

Definition 2.9.4 (eager-scope λ-term-graphs)**.** Let $G = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle$ be a λ-term-graph over $\Sigma^\lambda_{0_1,\mathsf{S}_j}$ for $j \in \{1,2\}$ with abstraction-prefix function $\mathrm{P} : V \to V^*$.

We say that G is an *eager-scope λ-term-graph*, or that G *is eager-scope* if:

$$\left.\begin{array}{l} \forall v, w \in V \;\; \forall p \in V^* \\ \quad \mathrm{P}(v) = p\, w \;\; \wedge \;\; v \notin \mathrm{V}(\mathsf{S}) \\ \quad\quad \Rightarrow \;\; \exists n \in \mathbb{N} \;\; \exists v_1, \ldots, v_n \in V \\ \quad\quad\quad\quad v \twoheadrightarrow v_1 \twoheadrightarrow \ldots \twoheadrightarrow v_n \twoheadrightarrow_0 w \;\wedge \\ \quad\quad\quad\quad v_n \in \mathrm{V}(\mathsf{0}) \;\; \wedge \;\; \forall i \in \{1, \ldots, n\} \;\; \mathrm{P}(v) \leq \mathrm{P}(v_i) \end{array}\right\} \tag{2.4}$$

In words: from every non-delimiter vertex v with non-empty abstraction-prefix $\mathrm{P}(v)$ ending with w there is a path to w via vertices with abstraction-prefixes that extend $\mathrm{P}(v)$ (a path within the scope of w) and via a variable vertex just before reaching w.

By ${}^{\mathrm{eag}}\mathcal{T}^{\lambda}_{1,j}$ we denote the subclass of $\mathcal{T}^{\lambda}_{1,j}$ of all eager-scope λ-term-graphs.

§ 2.9.5 (fully backlinked λ-term-graphs)**.** We will find that not only the class of eager λ-term-graphs is closed under functional bisimulation, but also a super-class thereof, called 'fully backlinked' λ-term-graphs.

Definition 2.9.6 (fully backlinked λ-term-graphs)**.** Let $G = \langle V, \mathrm{lab}, \mathrm{args}, r \rangle$ be a λ-term-graph over $\Sigma^{\lambda}_{0_1,\mathsf{S}_j}$ for $j \in \{1, 2\}$ with abstraction-prefix function $\mathrm{P} : V \to V^*$. We say that G is *fully backlinked* if:

$$\left.\begin{array}{l} \forall v, w \in V \;\; \forall p \in V^* \\ \quad \mathrm{P}(v) = p\, w \;\; \Rightarrow \;\; \exists n \in \mathbb{N} \;\; \exists v_1, \ldots, v_n \in V \\ \quad\quad\quad\quad v \twoheadrightarrow v_1 \twoheadrightarrow \ldots \twoheadrightarrow v_n \twoheadrightarrow w \\ \quad\quad\quad\quad \wedge \;\; \forall i \in \{1, \ldots, n\} \;\; \mathrm{P}(v) \leq \mathrm{P}(v_i) \end{array}\right\} \tag{2.5}$$

In words: from every vertices v with non-empty abstraction prefix $\mathrm{P}(v)$ that ends with w, there is a path from v to w via vertices with abstraction-prefixes that extend $\mathrm{P}(v)$ (a path within the scope of w).

By ${}^{\mathrm{fbl}}\mathcal{T}^{\lambda}_{1,j}$ we denote the subclass of $\mathcal{T}^{\lambda}_{1,j}$ that consists of all fully backlinked λ-term-graphs.

Proposition 2.9.7 (${}^{\mathrm{eag}}\mathcal{T}^{\lambda}_{1,2} \subseteq {}^{\mathrm{fbl}}\mathcal{T}^{\lambda}_{1,2}$)**.** Every eager-scope λ-term-graph over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ is fully backlinked.

Proof. Let G be an λ-term-graph over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ with abstraction-prefix function P and vertex set V. Suppose that G is an eager-scope λ-term-graph. This means that the condition (2.4) from definition 2.9.4 holds for G. Now note that, for

all $w, v \in V$ and $p \in V^*$, if $\mathrm{P}(w) = p\,v$ and $w \in \mathrm{V}(\mathsf{S})$ holds, then the correctness condition (S_1) on the abstraction-prefix function P entails $w \rightarrowtail_1 v$. It follows that G also satisfies the condition (2.5) from definition 2.9.6, and therefore, that G is fully backlinked. □

§ 2.9.8 (backlinks and fully backlinkedness)**.** The intuition for the property of a λ-term-graph G to be 'fully backlinked' is that, for every vertex w of G with a non-empty abstraction-prefix $\mathrm{P}(w) = p\,v$, it is possible to get back to the final abstraction vertex v in the abstraction-prefix of w by a directed path via vertices in the scope of v and via a last edge that is a backlink from a variable or a delimiter vertex to the abstraction vertex v. Therefore the presence of both sorts of backlink in λ-term-graphs is crucial for this concept. Indeed, at least backlinks for variables have to be present so that the property to be fully backlinked can make sense for a λ-term-graph: all λ-term-graphs with variable vertices but without variable backlinks (for example, consider a representation of the λ-term $\lambda x.\,x$ by such a term graph) are not fully backlinked.

§ 2.9.9 (eager-scope λ-term-graphs without variable backlinks)**.** The situation is different for the eager-scope property. While the presence of backlinks was used for the definition of 'eager-scope λ-term-graphs' in (2.4), the assumption that all variable vertices have backlinks is not essential there. In fact the condition (2.4) can be generalised to apply also to λ-term-graphs over $\Sigma^{\lambda}_{0_0,\mathsf{S}_1}$ and $\Sigma^{\lambda}_{0_0,\mathsf{S}_2}$ as follows:

$$\left.\begin{array}{l} \forall v, w \in V \;\; \forall p \in V^* \\ \quad \mathrm{P}(v) = p\,w \;\wedge\; v \notin \mathrm{V}(\mathsf{S}) \\ \qquad \Rightarrow \;\; \exists n \in \mathbb{N}\; \exists v_1, \ldots, v_n \in V \\ \qquad\qquad v \rightarrowtail v_1 \rightarrowtail \ldots \rightarrowtail v_n \\ \qquad\qquad \wedge \;\; v_n \in \mathrm{V}(\mathsf{0}) \;\wedge\; \mathrm{P}(v_n) = p\,w \\ \qquad\qquad \wedge \;\; \forall i \in \{1, \ldots, n\}\; \mathrm{P}(v) \leq \mathrm{P}(v_i) \end{array}\right\} \qquad (2.6)$$

For λ-term-graphs over $\Sigma^{\lambda}_{0_1,\mathsf{S}_1}$ and $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ the conditions (2.4) and (2.6) coincide: For the implication (2.6) ⇒ (2.4) note that the correctness condition $(\mathsf{0}_0)$ on the abstraction-prefix functions of these λ-term-graphs yields that the statements $v_n \in \mathrm{V}(\mathsf{0})$ and $\mathrm{P}(v_n) = \mathrm{P}(v) = p\,w$ imply that $v_n \rightarrowtail_0 w$. For the implication (2.4) ⇒ (2.6) observe that $v_n \in \mathrm{V}(\mathsf{0})$, $v_n \rightarrowtail_0 w$, and $\mathrm{P}(v_n) \leq p\,w$ implies $\mathrm{P}(v_n) = \mathrm{P}(v) = p\,w$ in view of the condition $(\mathsf{0}_0)$ and lemma 2.7.8.

Proposition 2.9.10. Functional bisimulation on λ-term-graphs in $\mathcal{T}^{\lambda}_{1,j}$ with $j \in \{1,2\}$ preserves and reflects the properties 'eager scope' and 'fully backlinked'.

Proof. We only show preservation under homomorphic image of the eager-scope property for λ-term-graphs, since preservation of the property 'fully backlinked' can be shown analogously and involves less technicalities. Also, reflection of these properties under functional bisimulation can be demonstrated similarly.

Let $j \in \{1,2\}$. Let $G_1 = \langle V_1, \mathrm{lab}_1, \mathrm{args}_1, r_1\rangle$ and $G_2 = \langle V_2, \mathrm{lab}_2, \mathrm{args}_2, r_2\rangle$ be λ-term-graphs over $\Sigma^{\lambda}_{0_1,\mathsf{S}_j}$ with (correct) abstraction-prefix functions P_1 and P_2, respectively. Let $h : V_1 \to V_2$ be a homomorphism from G_1 to G_2, and suppose that G_1 is eager-scope. We show that also G_2 is an eager-scope λ-term-graph.

For this, let $w' \in V_2$ such that $w' \notin \mathrm{V}(\mathsf{S})$ and $\mathrm{P}_2(w') = p'\,v'$ for some $v' \in V_2$ and $p' \in (V_2)^*$. Since h is surjective by proposition 2.2.12 (iii), there exists a vertex $w \in V_1$ such that $h(w) = w'$. Now note that, due to proposition 2.7.15, P_2 is the homomorphic image of P_1. It follows that $h^*(\mathrm{P}_1(w)) = \mathrm{P}_2(h(w)) = \mathrm{P}_2(w') = p'\,v'$. Hence there exist $v \in V_1$ and $p \in (V_1)^*$ such that $\mathrm{P}_1(w) = p\,v$ and $h^*(p) = p'$ and $h(v) = v'$. Since G_1 is an eager-scope λ-term-graph, there exists a path in G_1 of the form:

$$\pi :\ w = w_0 \rightarrowtail w_1 \rightarrowtail \ldots \rightarrowtail w_n \rightarrowtail_0 v$$

such that $w_n \in \mathrm{V}_1(\mathsf{0})$ and $\mathrm{P}_1(w_i) \geq p\,v$ for all $i \in \{0,\ldots,n\}$. As h is a homomorphism, it follows from proposition 2.2.12 (i) that π has an image $h(\pi)$ in G_2 of the form:

$$h(\pi) :\ w' = h(w) = h(w_0) \rightarrowtail' h(w_1) \rightarrowtail' \ldots \rightarrowtail' h(w_n) \rightarrowtail'_0 h(v) = v'$$

where $\rightarrowtail'$ is the directed-edge relation in G_2. Using again that h is a homomorphism, it follows that $h(w_n) \in \mathrm{V}_2(\mathsf{0})$. Due to the fact that P_2 is the homomorphic image of P_1 it follows that for all $i \in \{0,\ldots,n\}$ it holds: $\mathrm{P}_2(h(w_i)) = h^*(\mathrm{P}_1(w_i)) \geq h^*(p\,v) = h^*(p)\,h(v) = p'\,v'$. Hence we have shown that for $w'_i := h(w_i) \in V_2$ with $i \in \{0,\ldots,n\}$ it holds that $w' = w'_0 \rightarrowtail' w'_1 \rightarrowtail' \ldots \rightarrowtail' w'_n \rightarrowtail'_0 v'$ such that $w'_n \in \mathrm{V}_2(\mathsf{0})$ and $\mathrm{P}_2(w'_i) \geq p'\,w'$ for all $i \in \{0,\ldots,n\}$. In this way we have shown that also G_2 is eager-scope.

For showing that the eager-scope property is reflected by a homomorphism h from G_1 to G_2 the fact that paths in G_2 have pre-images under h in G (proposition 2.2.12 (ii)) can be used. □

§ 2.9.11 (generalised conditions for eager-scope and fully backlinked)**.** The following proposition states that the defining conditions for a λ-term-graph

to be eager-scope (definition 2.9.4) or fully backlinked (definition 2.9.6) can be generalised. The generalised condition for eager-scopedness requires that for every non-delimiter vertex v with $\mathrm{P}v = p\,w\,q$ there exists a path from v to w within the scope of w that only transits variable-vertex backlinks, but not delimiter-vertex backlinks. In the generalised condition for fully-backlinkedness the conclusion, in which the path may also proceed via delimiter-vertex backlinks, holds for all vertices.

Proposition 2.9.12 (generalised conditions for eager-scope and fully backlinked)**.** Let G be a λ-term-graph over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$, and let P be its abstraction-prefix function.

(i) G is an eager-scope λ-term-graph if and only if:

$$\left.\begin{array}{l} \forall v, w \in V \;\; \forall p, q \in V^* \\ \quad \mathrm{P}(v) = p\,w\,q \;\wedge\; v \notin \mathrm{V}(\mathsf{S}) \\ \qquad \Rightarrow \;\; \exists n \in \mathbb{N} \;\; \exists v_1, \ldots, v_n \in V \\ \qquad\qquad v \twoheadrightarrow v_1 \twoheadrightarrow \ldots \twoheadrightarrow v_n \twoheadrightarrow w \;\wedge \\ \qquad\qquad v_n \in \mathrm{V}(\mathsf{0}) \;\wedge\; \forall 0 \leq i \leq n \;\; p\,w \leq \mathrm{P}(v_i) \end{array}\right\} \tag{2.7}$$

(ii) The λ-term-graph G is fully backlinked if and only if:

$$\left.\begin{array}{l} \forall v, w \in V \;\; \forall p, q \in V^* \\ \quad \mathrm{P}(v) = p\,w\,q \;\Rightarrow\; \exists n \in \mathbb{N} \;\; \exists v_1, \ldots, v_n \in V \\ \qquad\qquad v \twoheadrightarrow v_1 \twoheadrightarrow \ldots \twoheadrightarrow v_n \twoheadrightarrow w \\ \qquad\qquad \wedge \;\; \forall 0 \leq i \leq n \;\; p\,w \leq \mathrm{P}(v_i) \end{array}\right\} \tag{2.8}$$

Proof. We will only prove the statement (ii), since statement (i) can be proven in much the same way.

Let G be an arbitrary λ-term-graph over $\Sigma^{\lambda}_{0_1,\mathsf{S}_j}$ with $j \in \{1, 2\}$. If (2.8) holds for G, then so does its special case (2.5), and it follows that G is fully backlinked.

For showing the converse we assume that G is fully backlinked. Then (2.5) holds. We show (2.8), which is universally quantified over $q \in V^*$, by induction on the length $|q|$ of q.

In the base case we have $q = \epsilon$, and the statement to be establish follows immediately from (2.5).

For the induction step, let $p, q \in V^*$ and $v, w \in V$ be such that $\mathrm{P}(v) = p\,w\,q$ with $q = w'\,q_0$ for $w' \in V$ and $q_0 \in V^*$. Due to $|q_0| < |q|$, the induction hypothesis

can be applied to $\mathrm{P}(v) = (p\,w)\,w'\,q_0$, yielding a path $v \rightarrowtail^* v' \rightarrowtail w'$, for some $v' \in V$, that in its part between v and v' visits vertices with $p\,w\,w'$ as prefix of their abstraction prefixes. Due to $p\,w\,w' \leq \mathrm{P}(v')$ it follows that $\mathrm{P}(w') = p\,w$ by lemma 2.7.8, (i) (which entails that $p\,w\,w' = \mathrm{P}(v')$ and that $v' \rightarrowtail w'$ must be a backlink). From this the condition (2.5) yields a path $w' \rightarrowtail^* v'' \rightarrowtail w$, for some $v'' \in V$, that on the way from w' to v'' visits vertices with $p\,w$ as prefix of their abstraction prefix. Combining these two paths, we obtain a path $v \rightarrowtail^* v' \rightarrowtail w' \rightarrowtail^* v'' \rightarrowtail w$ that before it reaches v'' visits only vertices with $p\,w$ as prefix of their abstraction prefixes. This establishes the statement to show for the induction step. □

§ 2.9.13. In order to show that eager-scope and fully backlinked λ-term-graphs are closed under functional bisimulation, we first establish the following lemma: if two vertices of such a λ-term-graph have the same image under a homomorphism h, then so do their abstraction prefixes.

Lemma 2.9.14. Let $G \in \mathcal{T}^{\lambda}_{1,2}$ be a fully backlinked λ-term-graph with vertex set V and abstraction-prefix function P. Let $G' \in \mathcal{T}_{1,2}$ be a term graph over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ such that $G \rightrightarrows_h G'$ for a homomorphism h. Then it holds:

$$\forall v_1, v_2 \in V \;\; h(v_1) = h(v_2) \;\Rightarrow\; h^*(\mathrm{P}(v_1)) = h^*(\mathrm{P}(v_2)) \tag{2.9}$$

where h^* is the homomorphic extension of h to words over V.

To prove this lemma, as a stepping stone, we first prove the following technical lemma. To this end we use the generalised condition for fully backlinkedness from proposition 2.9.12 to relate the abstraction-prefixes of vertices on paths departing from v_1 and v_2.

Lemma 2.9.15. Let $G = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle \in \mathcal{T}^{\lambda}_{1,j}$ for $j \in \{1,2\}$ with abstraction-prefix function P. Let $G' = \langle V', \mathrm{lab}', \mathrm{args}', r'\rangle \in \mathcal{T}_{1,j}$ (i.e. G' is not necessarily a λ-term-graph) such that $G \rightrightarrows_h G'$ for a homomorphism h. Let $v_1, v_2 \in V$ be such that $h(v_1) = h(v_2)$.

Suppose that $\mathrm{P}(v_1) = p_1\,w_1\,q_1$ with $w_1 \in V$ and $p_1, q_1 \in V^*$ and that π_1 is a path from v_1 to w_1 in G of the form

$$\pi_1 : v_1 = v_{1,0} \rightarrowtail v_{1,1} \rightarrowtail \ldots \rightarrowtail v_{1,n-1} \rightarrowtail v_{1,n} = w_1$$

such that furthermore $\mathrm{P}(v_{1,n-1}) = p_1\,w_1$ holds.

Then there are $w_2 \in V$ and $p_2, q_2 \in V^*$ with $|q_2| = |q_1|$ such that $\mathrm{P}(v_2) = p_2\, w_2\, q_2$, and a path π_2 in G from v_2 to w_1 of the form

$$\pi_2 : v_2 = v_{2,0} \rightarrowtail v_{2,1} \rightarrowtail \ldots \rightarrowtail v_{2,n-1} \rightarrowtail v_{2,n} = w_2$$

such that $\mathrm{P}(v_{2,n-1}) = p_2\, w_2$, and that $h(v_{1,j}) = h(v_{2,j})$ for all $j \in \{0, \ldots, n\}$, and in particular $h(w_1) = h(w_2)$.

Proof. By lemma 2.7.8 (i), from $\mathrm{P}(v_{1,n-1}) = p_1\, w_1$ it follows that $\mathrm{P}(w_1) = p_1$. Hence the final transition in π_1 must be either via the backlink of a variable vertex or a delimiter vertex and is therefore either of the form $v_{1,n-1} \overset{0}{\rightarrowtail}_0 v_{1,n}$ or of the form $v_{1,n-1} \overset{S}{\rightarrowtail}_1 v_{1,n}$.

Furthermore by proposition 2.2.12 (i), it follows that there is a path $h(\pi_1)$ in G' from $h(v_1)$ to $h(w_1)$ of the form:

$$h(\pi_1) : h(v_1) = h(v_{1,0}) \rightarrowtail h(v_{1,1}) \rightarrowtail \ldots \rightarrowtail h(v_{1,n-1}) \rightarrowtail h(v_{1,n}) = h(w_1)$$

From this path, proposition 2.2.12 (ii), yields a path π_2 in G from v_2 of the form:

$$\pi_2 : v_2 = v_{2,0} \rightarrowtail v_{2,1} \rightarrowtail \ldots \rightarrowtail v_{2,n-1} \rightarrowtail v_{2,n} = w_2$$

such that $h(\pi_2) = h(\pi_1)$. Therefore it holds $h(v_{1,j}) = h(v_{2,j})$ for all $j \in \{1, \ldots, n\}$, and in particular $h(w_1) = h(w_2)$. Since h is a homomorphism also $\mathrm{lab}(v_{1,j}) = \mathrm{lab}(v_{2,j})$ follows for all $j \in \{0, \ldots, n\}$. Therefore, and due to lemma 2.7.8 (iv), the abstraction-prefix function P quantitatively behaves the same when stepping through π_2 as when stepping through π_1. It follows that $\mathrm{P}(v_2) = p_2\, w\, q_2$, $\mathrm{P}(v_{2,n-1}) = p_2\, w$, and $\mathrm{P}(w_2) = p_2$ for some $w \in V$, and $p_2, q_2 \in V^*$ with $|q_2| = |q_1|$. Due to $\mathrm{lab}(v_{2,n-1}) = \mathrm{lab}(v_{1,n-1})$, and since the final transition $v_{1,n-1} \rightarrowtail w_1$ in π_1 is one along a backlink of a variable or delimiter vertex, this also holds for the final transition $v_{2,n-1} \rightarrowtail w_2$ in π_2. Then it follows by the correctness condition $(\mathsf{0}_1)$ or (S_1), respectively, that $w = w_2$, and therefore that $\mathrm{P}(v_2) = p_2\, w_2\, q_2$, and $\mathrm{P}(v_{2,n-1}) = p_2\, w_2$. □

Relying on this lemma, we can now give a rather straightforward proof of lemma 2.9.14.

Proof of lemma 2.9.14. Let G, G' be as assumed in the lemma, and let h be a homomorphism that witnesses $G \mathrel{\dot{\geq}}_h G'$. We first show:

$$\left.\begin{array}{l} \forall v_1, v_2 \in V\ \forall w_1 \in V\ \forall p_1, q_1 \in V^* \\ \quad h(v_1) = h(v_2)\ \wedge\ \mathrm{P}(v_1) = p_1\, w_1\, q_1\ \Longrightarrow \\ \qquad \exists w_2 \in V\ \exists p_2, q_2 \in V^*\ \mathrm{P}(v_2) = p_2\, w_2\, q_2\ \wedge \\ \qquad\qquad h(w_1) = h(w_2)\ \wedge\ |q_2| = |q_1| \end{array}\right\} \tag{2.10}$$

For this, let $v_1, v_2 \in V$ be such that $h(v_1) = h(v_2)$ and $\mathrm{P}(v_1) = p_1\, w_1\, q_1$ for $p_1, q_1 \in V^*$ and $w_1 \in V$. Now since G is fully backlinked, there is a path $\pi_1 : v_1 = v_{1,0} \rightarrowtail^* v_{1,n-1} \rightarrowtail v_{1,n} = w_1$ in G with $\mathrm{P}(v_{1,n-1}) = p_1\, w_1$. Then lemma 2.9.15 yields the existence of $p_1, q_1 \in V^*$ and $w_1 \in V$ with $|q_2| = |q_1|$ such that $h(w_1) = h(w_2)$, and $\mathrm{P}(v_2) = p_2\, w_2\, q_2$. This establishes (2.10).

As an direct consequence of (2.10) we obtain:

$$\forall v_1, v_2 \in V\ \forall r_1 \in V^*\ h(v_1) = h(v_2)\ \wedge\ \mathrm{P}(v_1) = r_1\ \Longrightarrow \\ \exists s_2, r_2 \in V^*\ \mathrm{P}(v_2) = s_2\, r_2\ \wedge\ h^*(r_1) = h^*(r_2)$$

But since in this statement the roles of v_1 and v_2 can be exchanged, it follows that $h(v_1) = h(v_2)$ always entails $|\mathrm{P}(v_1)| = |\mathrm{P}(v_2)|$, and hence $h^*(\mathrm{P}(v_1)) = h^*(\mathrm{P}(v_2))$. This establishes (2.9). □

Lemma 2.9.14 is the crucial stepping stone for the proof of the main theorem of this chapter.

Theorem 2.9.16 (preservation of the properties 'eager scope' and 'fully backlinked' under functional bisimulation on $\mathcal{T}^{\lambda}_{1,2}$)**.** Let G be a λ-term-graph over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ and suppose that h is a homomorphism from G to a term graph $G' \in \mathcal{T}_{1,2}$. Then the following two statements hold:

(i) If G is fully backlinked, then also G' is a λ-term-graph, which is fully backlinked.

(ii) If G is eager-scope, then also G' is a λ-term-graph, which is eager scope.

Proof. Let $G = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle \in \mathcal{T}^{\lambda}_{1,2}$ and $G' = \langle V', \mathrm{lab}', \mathrm{args}', r'\rangle \in \mathcal{T}_{1,2}$, and h a homomorphism from G to G'.

For showing statement (i) of the theorem, we assume that G is fully backlinked. On G' we define the following abstraction-prefix function:

$$\begin{aligned} \mathrm{P}' \ :\ & V' \to (V')^*, \\ & w' \mapsto h^*(\mathrm{P}(w))\ \text{ for some } w \in V \text{ with } h(w) = w' \end{aligned} \tag{2.11}$$

This function is well-defined because: for every $w' \in V'$ there exists a $w \in V$ with $w' = h(w)$, due to proposition 2.2.12 (iii); and due to lemma 2.9.14 the $\mathrm{P}'(w')$ is the same for all $w \in V$ with $w' = h(w)$.

For the thus defined abstraction-prefix function P' it holds:

$$\forall w \in V\quad h^*(\mathrm{P}(w)) = \mathrm{P}'(h(w))$$

and hence P' is the homomorphic image of P under h in the sense of definition 2.7.12.

It remains to show that P' is a correct abstraction-prefix function for G', and that G' is fully backlinked. Both properties follow from statements established earlier by using that G' and P' are the homomorphic images under h of G and P, respectively: That P' is a correct abstraction-prefix function for G' follows from lemma 2.7.16 (ii) which entails that the homomorphic image of a correct abstraction-prefix function is correct. And that G' is fully backlinked follows from proposition 2.9.10, which states that the homomorphic image of a fully backlinked λ-term-graph is again fully backlinked.

In this way we have established statement (i) of the theorem.

For showing statement (ii), let G be an eager-scope λ-term-graph over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$. By proposition 2.9.7 it follows that G is also fully backlinked. Therefore the just established statement (i) of the theorem is applicable, and it yields that G' is a λ-term-graph. Since by proposition 2.9.10 also the eager-scope property is preserved by homomorphism, it follows that G' is eager-scope, too. This establishes statement (ii) of the theorem. □

Corollary 2.9.17. ${}^{\mathrm{fbl}}\mathcal{T}^{\lambda}_{1,2}$ and ${}^{\mathrm{eag}}\mathcal{T}^{\lambda}_{1,2}$ are closed under functional bisimulation.

§ 2.9.18 (backlinks are required)**.** Note that statements analogous to theorem 2.9.16 and corollary 2.9.17 do not hold for λ-term-graphs over $\Sigma^{\lambda}_{0_1,\mathsf{S}_1}$: The classes ${}^{\mathrm{fbl}}\mathcal{T}^{\lambda}_{1,1}$ and ${}^{\mathrm{eag}}\mathcal{T}^{\lambda}_{1,1}$ are not closed under functional bisimulation. This is witnessed by the counterexample in the proof of proposition 2.8.4 (iii), which maps an eager-scope (and hence fully-backlinked) λ-term-graph to a term graph that is not a λ-term-graph.

Similarly the statements of theorem 2.9.16 and corollary 2.9.17 do not carry over to λ-term-graphs over $\Sigma^{\lambda}_{0_0,\mathsf{S}_1}$ and $\Sigma^{\lambda}_{0_0,\mathsf{S}_2}$ that are eager-scope in the sense of § 2.9.9. This is witnessed by the eager-scope term graphs in the proof of proposition 2.8.4 (i) and (ii).

Another direct consequence of theorem 2.9.16 is the following corollary. Recall the notations from definition 2.2.9 and Notation 2.2.16.

Corollary 2.9.19. The following statements holds:

(i) For every fully backlinked λ-term-graph G over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ it holds:

$$(\boldsymbol{G}\twoheadrightarrow) = (\boldsymbol{G}\twoheadrightarrow)^{{}^{\mathrm{fbl}}\mathcal{T}^{\lambda}_{1,2}} = (\boldsymbol{G}\twoheadrightarrow)^{\mathcal{T}^{\lambda}_{1,2}}$$

(ii) For every eager-scope λ-term-graph G over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ it holds:

$$(\boldsymbol{G}\twoheadrightarrow) = (\boldsymbol{G}\twoheadrightarrow)^{{}^{\mathrm{eag}}\mathcal{T}^{\lambda}_{1,2}} = (\boldsymbol{G}\twoheadrightarrow)^{\mathcal{T}^{\lambda}_{1,2}}$$

This corollary will be central to proving in section 2.10 the complete-lattice property for $\twoheadrightarrow$-successors of ($\sim$-equivalence classes of) λ-term-graphs over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ and λ-ap-ho-term-graphs over $\Sigma^{\lambda}_{0_1}$.

Remark 2.9.20 (an alternative remedy)**.** To recapitulate, we identify subclasses of $\mathcal{T}^{\lambda}_{1,2}$ that are closed under functional bisimulation. These subclasses are ${}^{\mathrm{fbl}}\mathcal{T}^{\lambda}_{1,2}$ and the subclass ${}^{\mathrm{eag}}\mathcal{T}^{\lambda}_{1,2}$ thereof. The reason why some λ-term-graphs are not fully backlinked is revealed by the counterexample of proposition 2.8.4 (iv): in the scopes of the two topmost λ-abstractions in G_1, the two subgraphs representing $\lambda x.\, x$ are 'dangling' since the scopes of the topmost λ-abstractions are not closed; therefore the mentioned subgraphs can be shared in the homomorphic image G_0, thus leading to overlapping scopes.

In the original paper [24, 7.12], we explore an approach to mend this blemish of $\mathcal{T}^{\lambda}_{1,2}$ by changing its definition. We ensure that *all* λ-term-graphs are fully backlinked. This is facilitated by allowing trailing S-vertices as successors of $\mathsf{0}$-vertices. Thereby scopes that remain open at a variable occurrence are still closed afterwards. This solution is omitted in this thesis.

2.10 Transfer of the complete-lattice property to λ-ho-term-graphs

§ 2.10.1 (overview)**.** In this section we establish that sets of $\twoheadrightarrow$-successors of (the isomorphism equivalence class of) a given λ-term-graph over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ form a complete lattice under the sharing order. For this we use the fact that this is the case for first-order term graphs in general (see proposition 2.2.18), and we apply the results developed so far. Subsequently we transfer this complete-lattice property to the higher-order λ-ap-ho-term-graphs over $\Sigma^{\lambda}_{0_1}$ via the correspondences established in section 2.7.

The following proposition is a specialisation proposition 2.2.18.

Proposition 2.10.2 (complete-lattice property of $\mathcal{T}_{i,j}$)**.** $\langle(\boldsymbol{G}\twoheadrightarrow), \twoheadrightarrow\rangle$ is a complete lattice for every term graph G over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ with $i \in \{0,1\}$ and $j \in \{1,2\}$.

Combining this proposition with the result from section 2.9 that the classes of fully backlinked and eager-scope λ-term-graphs over $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ are closed under functional bisimulation (corollary 2.9.17) and with corollary 2.9.19 we obtain the following theorem.

Theorem 2.10.3 (complete-lattice property of ${}^{\mathrm{fbl}}\mathcal{T}^{\lambda}_{1,2}$ and ${}^{\mathrm{eag}}\mathcal{T}^{\lambda}_{1,2}$)**.** $\langle (G\twoheadrightarrow)^{\mathcal{T}^{\lambda}_{1,2}}, \twoheadrightarrow\rangle$ is a complete lattice for all $G \in {}^{\mathrm{fbl}}\mathcal{T}^{\lambda}_{1,2} \cup {}^{\mathrm{eag}}\mathcal{T}^{\lambda}_{1,2}$.

§ 2.10.4 (transfer to eager-scope higher-order term graphs)**.** Considering this theorem we cannot expect all higher-order term graphs to form a complete lattice, but only those subclasses that correspond to fully backlinked or eager-scope λ-term-graphs. However, the 'fully-backlinked' property does not have a natural equivalent on the higher-order term graphs, since it is based on the existence of paths that depend on backlinks departing from delimiter vertices, which the higher-order term graphs do not have. The eager-scope property on the other hand has an obvious counterpart in the higher-order term graphs. Its definition below is analogous to the eager-scope property of λ-term-graphs as described in § 2.9.9.

Definition 2.10.5 (eager-scope λ-ap-ho-term-graphs)**.** Let $i \in \{0,1\}$ and let $\mathcal{G} = \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{P}\rangle$ be a λ-ap-ho-term-graph over $\Sigma^{\lambda}_{0_1}$. We say that $\mathcal{G}$ is an *eager-scope λ-term-graph*, or that $\mathcal{G}$ is *eager-scope*, if:

$$\begin{aligned}
&\forall w, v \in V\ \forall p \in V^*\ \ \mathrm{P}(w) = p\,v\ \Rightarrow\\
&\quad \exists n \in \mathbb{N}\ \exists w_1, \ldots, w_n \in V\\
&\qquad w \twoheadrightarrow w_1 \twoheadrightarrow \ldots \twoheadrightarrow w_n\ \wedge\ w_n \in \mathrm{V}(0)\ \wedge\ \mathrm{P}(w_n) = p\,v\\
&\qquad \wedge\ \forall i \in \{1, \ldots, n-1\}\ \ p\,v \leq \mathrm{P}(w_i)
\end{aligned}$$

Or in other words, if for every vertex w of G with a non-empty abstraction-prefix $\mathrm{P}(w)$ ending with v there exists a path from w via vertices with abstraction-prefixes that start with $\mathrm{P}(w)$ (i.e. a path within the scope of v) to a variable vertex w_n that is bound by the abstraction vertex v. (Note that if $k = 1$, then v is directly reachable from w_n via its backlink.)

By ${}^{\mathrm{eag}}\mathcal{H}^{(\lambda)}_1$ we denote the subclass of $\mathcal{H}^{(\lambda)}_1$ that consists of all eager-scope λ-ap-ho-term-graphs.

Proposition 2.10.6 (uniqueness of eager-scope λ-ap-ho-term-graphs)**.** Let $\mathcal{G}_i = \langle V, \mathrm{lab}, \mathrm{args}, r, \mathrm{P}_i\rangle$ with $i \in \{1,2\}$ be λ-ap-ho-term-graphs with the same underlying term graph. If $\mathcal{G}_1$ is eager-scope, then $|\mathrm{P}_1(w)| \leq |\mathrm{P}_2(w))|$ for all $w \in V$. If, in addition, also $\mathcal{G}_2$ is eager-scope, then $\mathrm{P}_1 = \mathrm{P}_2$. Hence eager-scope λ-ap-ho-term-graphs over the same underlying term graph are unique.

Proof sketch. If a λ-ap-ho-term-graph is not eager-scope, then it contains a vertex w with abstraction-prefix $v_1 \dots v_n$ from which v_n is only reachable (if at all) by leaving the scope of v_n. It can be shown that in this case another abstraction-prefix function with shorter prefixes exists in which v_n does not occur in the prefix of w. □

As stated in the proposition below, the eager-scope property for λ-ap-ho-term-graphs and for λ-term-graphs correspond to each other via the correspondence mappings from proposition 2.7.18 and proposition 2.7.19.

Proposition 2.10.7 ($\mathrm{G}_{i,j}$ and $\mathcal{G}_{i,j}$ preserve and reflect the eager-scope property)**.** Let $i \in \{0,1\}$, and $j \in \{1,2\}$. The correspondences $\mathrm{G}_{i,j}$ and $\mathcal{G}_{i,j}$ between preserve and reflect the eager-scope property i.e., for all $\mathcal{G} \in \mathcal{H}_i^{(\lambda)}$, and for all $G \in \mathcal{T}_{i,j}^{\lambda}$ it holds:

$$\begin{aligned} \mathcal{G} \text{ is eager-scope} &\;\Leftrightarrow\; \mathrm{G}_{i,j}(\mathcal{G}) \text{ is eager-scope} \\ \mathcal{G}_{i,j}(G) \text{ is eager-scope} &\;\Leftrightarrow\; G \text{ is eager-scope} \end{aligned}$$

Consequently, the restriction of the domains of $\mathrm{G}_{i,j}$ and $\mathcal{G}_{i,j}$ to eager-scope term graphs – and also of $\mathbf{G}_{i,j}$ and $\boldsymbol{\mathcal{G}}_{i,j}$ – are functions of following types:

$$\begin{aligned} \mathrm{G}_{i,j}|_{^{\mathrm{eag}}\mathcal{H}_i^{(\lambda)}} &: {}^{\mathrm{eag}}\mathcal{H}_i^{(\lambda)} \to {}^{\mathrm{eag}}\mathcal{T}_{i,j}^{\lambda} & \mathcal{G}_{i,j}|_{^{\mathrm{eag}}\mathcal{T}_{i,j}^{\lambda}} &: {}^{\mathrm{eag}}\mathcal{T}_{i,j}^{\lambda} \to {}^{\mathrm{eag}}\mathcal{H}_i^{(\lambda)} \\ \mathbf{G}_{i,j}|_{^{\mathrm{eag}}\boldsymbol{\mathcal{H}}_i^{(\boldsymbol{\lambda})}} &: {}^{\mathrm{eag}}\boldsymbol{\mathcal{H}}_i^{(\boldsymbol{\lambda})} \to {}^{\mathrm{eag}}\boldsymbol{\mathcal{T}}_{i,j}^{(\boldsymbol{\lambda})} & \boldsymbol{\mathcal{G}}_{i,j}|_{^{\mathrm{eag}}\boldsymbol{\mathcal{T}}_{i,j}^{(\boldsymbol{\lambda})}} &: {}^{\mathrm{eag}}\boldsymbol{\mathcal{T}}_{i,j}^{(\boldsymbol{\lambda})} \to {}^{\mathrm{eag}}\boldsymbol{\mathcal{H}}_i^{(\boldsymbol{\lambda})} \end{aligned}$$

Furthermore, we can specialise theorem 2.7.20 concerning the correspondences on the isomorphism equivalence classes (in particular (i) and (iii)) as follows. Recall Notation 2.2.16.

Lemma 2.10.8. Let $i \in \{0,1\}$, and $j \in \{1,2\}$ and let $\mathcal{G} \in \mathcal{H}_i^{(\lambda)}$ be a λ-ap-ho-term-graph. Then

$$\begin{aligned} \mathbf{G}_{i,j}|_{(\boldsymbol{\mathcal{G}}\twoheadrightarrow)} &: (\boldsymbol{\mathcal{G}} \twoheadrightarrow) \to (\mathbf{G}_{i,j}(\boldsymbol{\mathcal{G}})\twoheadrightarrow)^{\mathcal{T}_{i,j}^{\lambda}} \\ \boldsymbol{\mathcal{G}}_{i,j}|_{(\mathbf{G}_{i,j}(\boldsymbol{\mathcal{G}})\twoheadrightarrow)^{\mathcal{T}_{i,j}^{\lambda}}} &: (\mathbf{G}_{i,j}(\boldsymbol{\mathcal{G}})\twoheadrightarrow)^{\mathcal{T}_{i,j}^{\lambda}} \to (\boldsymbol{\mathcal{G}} \twoheadrightarrow) \end{aligned}$$

and it holds that $\boldsymbol{\mathcal{G}}_{i,j}|_{(\mathbf{G}_{i,j}(\boldsymbol{\mathcal{G}})\twoheadrightarrow)^{\mathcal{T}_{i,j}^{\lambda}}}$ is a left-inverse of $\mathbf{G}_{i,j}|_{(\boldsymbol{\mathcal{G}}\twoheadrightarrow)}$:

$$\boldsymbol{\mathcal{G}}_{i,j}|_{(\mathbf{G}_{i,j}(\boldsymbol{\mathcal{G}})\twoheadrightarrow)^{\mathcal{T}_{i,j}^{\lambda}}} \circ \mathbf{G}_{i,j}|_{(\boldsymbol{\mathcal{G}}\twoheadrightarrow)} = \mathrm{id}_{(\boldsymbol{\mathcal{G}}\twoheadrightarrow)}$$

The final tool for transfering the complete-lattice property to higher-order term graphs is the following general lemma about partial orders. It states that the complete-lattice property of partial orders is reflected by order homomorphisms with left-inverses.

Lemma 2.10.9 (reflection of the complete-lattice property under order homomorphisms with left-inverses)**.** Let $\langle A, \leq_A \rangle$ and $\langle B, \leq_B \rangle$ be partial orders. Suppose that $h : A \to B$, and $i : B \to A$ are order homomorphisms such that i is a left-inverse of h, i.e.: $i \circ h = \mathrm{id}_A$. Then if $\langle B, \leq_B \rangle$ is a complete lattice, then so is $\langle A, \leq_A \rangle$.

Proof. Let $S \subseteq A$ be arbitrary. We have to show the existence of a least upper bound (l.u.b.) $\bigsqcup S$ and a greatest lower bound (g.l.b.) $\bigsqcap S$ of S in $\langle A, \leq_A \rangle$. We only show the existence of the l.u.b., because the argument for the g.l.b. is analogous. Since $\langle B, \leq_B \rangle$ is a complete lattice, we know that a l.u.b. $\bigsqcup h(S)$ exists. We will show that $\bigsqcup S = i(\bigsqcup h(S))$.

$i(\bigsqcup h(S))$ is an upper bound of S: Let $s \in S$ be arbitrary. Then $h(s) \in h(S)$, and $h(s) \leq_B \bigsqcup h(S)$. Since i is an order homomorphism, and a left-inverse of h, it follows that $s = i(h(s)) \leq_A i(\bigsqcup h(S))$.

$i(\bigsqcup h(S))$ is less or equal to all upper bounds of S: Let $u \in A$ be an upper bound of S. As h is an order homomorphism, it follows that $h(u)$ is an upper bound of $h(S)$. Consequently, $\bigsqcup h(S) \leq_B h(u)$. Again, since i is an order homomorphism, and a left-inverse of h, this entails $i(\bigsqcup h(S)) \leq_A i(h(u)) = u$. □

Now we can formulate and prove the main result of this section.

Theorem 2.10.10. For every $\mathcal{G} \in {}^{\mathrm{eag}}\mathcal{H}_1^\lambda$ it holds that $\langle (\boldsymbol{\mathcal{G}} \twoheadrightarrow), \twoheadrightarrow \rangle$ is a complete lattice.

Proof. Let $\mathcal{G} \in {}^{\mathrm{eag}}\mathcal{H}_1^\lambda$ Then by proposition 2.10.7 it follows that $\mathrm{G}_{1,2}(\mathcal{G}) \in {}^{\mathrm{eag}}\mathcal{T}_{1,2}^\lambda$ and also that $\mathbf{G}_{1,2}(\boldsymbol{\mathcal{G}}) \in {}^{\mathrm{eag}}\boldsymbol{\mathcal{T}}_{1,2}^{(\boldsymbol{\lambda})}$. Now theorem 2.10.3 yields that $(\mathbf{G}_{1,2}(\boldsymbol{\mathcal{G}}){\twoheadrightarrow})^{\mathcal{T}_{1,2}^\lambda}$ w.r.t. $\twoheadrightarrow$ is a complete lattice. Furthermore note that $\boldsymbol{\mathcal{G}}_{1,2}\big|_{(\mathbf{G}_{1,2}(\boldsymbol{\mathcal{G}}){\twoheadrightarrow})^{\mathcal{T}_{1,2}^\lambda}}$ is a left-inverse of $\mathbf{G}_{1,2}\big|_{(\boldsymbol{\mathcal{G}}\twoheadrightarrow)}$ due to lemma 2.10.8. Hence lemma 2.10.9 can be used to show that the complete-lattice property of $(\mathbf{G}_{1,2}(\boldsymbol{\mathcal{G}}){\twoheadrightarrow})^{\mathcal{T}_{1,2}^\lambda}$ w.r.t. $\twoheadrightarrow$ is reflected by $\mathbf{G}_{1,2}\big|_{(\boldsymbol{\mathcal{G}}\twoheadrightarrow)}$, yielding that $\langle (\boldsymbol{\mathcal{G}} \twoheadrightarrow), \twoheadrightarrow \rangle$ is a complete lattice. □

2.11 Summary

§ 2.11.1 (summary)**.** We defined higher-order term graph representations for strongly regular λ-terms:

- λ-ho-term-graphs $\mathcal{H}_i^{\lambda}$, an adaptation of Blom's 'higher-order term graphs' [8], which possess a scope function that maps every abstraction vertex v to the set of vertices that are in the scope of v.
- λ-ap-ho-term-graphs $\mathcal{H}_i^{(\lambda)}$, which instead of a scope function carry an abstraction-prefix function that assigns to every vertex information about the scoping structure. Abstraction prefixes are closely related to the notion of 'generated subterms' for λ-terms (definition 1.4.39). The correctness conditions here are simpler and more intuitive than for λ-ho-term-graphs.

These classes are defined for $i \in \{0, 1\}$, according to whether variable occurrences have backlinks to abstractions or not. Our main statements about these classes are:

- a bijective correspondence between $\mathcal{H}_i^{\lambda}$ and $\mathcal{H}_i^{(\lambda)}$ via mappings A_i and B_i that preserve and reflect the sharing order (theorem 2.5.11);
- the naive approach to implementing functional bisimulation on theses classes (ignoring all scoping information and using only the underlying first-order term graphs) fails (proposition 2.6.6).

The latter was the motivation to consider first-order term graphs with scope delimiters:

- λ-term-graphs $\mathcal{T}_{i,j}^{\lambda}$ (with $i \in \{0, 1\}$ for whether there are variable backlinks and $j \in \{1, 2\}$ for whether there are delimiter backlinks), which are plain first-order term graphs, but which require as a correctness condition the existence of an abstraction-prefix function.

The most important results linking these classes with λ-ap-ho-term-graphs are:

- an 'almost bijective' correspondence between $\mathcal{H}_i^{(\lambda)}$ via mappings $\mathrm{G}_{i,j}$ and $\mathcal{G}_{i,j}$ that preserve and reflect the sharing order (theorem 2.7.20);
- the subclass ${}^{\mathrm{eag}}\mathcal{T}_{1,2}^{\lambda}$ of eager-scope λ-term-graphs is closed under functional bisimulation (corollary 2.9.17).

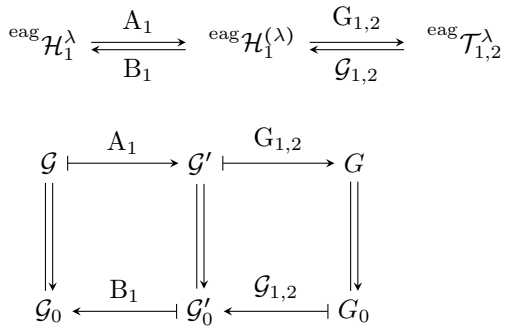

Figure 2.6. The correspondences (top) permit an implementation of functional bisimulation on higher-order term graphs using functional bisimulation on first-order term graphs (bottom).

The correspondences together with the closedness result allow us to handle functional bisimulation between eager-scope higher-order term graphs in a straightforward manner by implementing them via functional bisimulation between first-order term graphs as shown in fig. 2.6.

§ 2.11.2 (outlook: maximal sharing)**.** The findings from this chapter are a toehold for the concept of maximal sharing, which we develop in the following chapter. In particular the complete-property of the classes ${}^{\text{eag}}\mathcal{H}_1^\lambda$, ${}^{\text{eag}}\mathcal{T}_{1,2}^\lambda$, and $\mathcal{T}_{1,2}$ implies that for every graph in these classes there is a unique maximally compact version of that element – its bisimulation collapse – which we call its maximally shared form. For an eager λ-ho-term-graph $\mathcal{G}$ the maximally shared form can be computed as:

$$|\!\downarrow^{\text{eag}}_{\mathcal{H}_1^\lambda}(\mathcal{G}) = (\mathrm{B}_1 \circ \mathcal{G}_{1,2} \circ |\!\downarrow_{\mathcal{T}_{1,2}} \circ \mathrm{G}_{1,2} \circ \mathrm{A}_1)(\mathcal{G})$$

For implementing $|\!\downarrow_{\mathcal{T}_{1,2}}$ fast algorithms are available.

Analogous to remark 2.9.20 this can be generalised to term graphs without eager scope-closure. For our intent of getting a grip on maximal sharing in $\lambda_{\mathsf{letrec}}$, however, only eager scope-closure is practically relevant, because it facilitates the highest degree of sharing.

Chapter 3

Maximal Sharing in $\lambda_{\mathsf{letrec}}$

3.1 Overview

§ 3.1.0 (teaser)**.** Can we transform $\lambda f.\,\mathsf{let}\ r = f\ (f\ r)\ \mathsf{in}\ r$ into the more efficient form $\lambda f.\,\mathsf{let}\ r = f\ r\ \mathsf{in}\ r$ and can we prove that these terms are equivalent?

§ 3.1.1 (subject matter)**.** Increasing sharing in programs is generally desirable. It results in more compact code and avoids duplication of reduction work at run-time, thereby speeding up execution. We show how a maximal degree of sharing can be obtained for programs expressed as terms in the λ-calculus with letrec. We introduce a notion of 'maximal compactness' for $\lambda_{\mathsf{letrec}}$-terms among all terms with the same unfolding. We translate $\lambda_{\mathsf{letrec}}$-terms into a term graph representation which respects the unfolding semantics in the sense that bisimilarity preserves and reflects unfolding equivalence on the $\lambda_{\mathsf{letrec}}$-terms. Compactness of the term graphs can then be compared and increased by means of functional bisimulation.

§ 3.1.2 (methods and formalisms)**.** We lean heavily on the results from the previous chapters. While in chapter 1 we completely focused on the λ-terms expressed by $\lambda_{\mathsf{letrec}}$-terms, in chapter 2 we exclusively investigated possible suitable graph representations for $\lambda_{\mathsf{letrec}}$-terms. In this chapter we put things together and connect the graph representations to the unfolding semantics. We provide a translation into higher-order and first-order term graphs and we show that it respects the unfolding semantics in the sense that bisimulation preserves and reflects the unfolding semantics.

§ 3.1.3 (results)**.** We obtain practical and efficient methods for the following two problems: transforming a $\lambda_{\mathsf{letrec}}$-term into a maximally compact form; and deciding whether two $\lambda_{\mathsf{letrec}}$-terms have the same unfolding.

The transformation of a $\lambda_{\mathsf{letrec}}$-term L into maximally compact form L_0 proceeds in three steps:

(i) translate L into a term graph $G = [\![L]\!]$;

(ii) compute the maximally shared form of G as its bisimulation collapse G_0;

(iii) read back a $\lambda_{\mathsf{letrec}}$-term L_0 from G_0 with the property $[\![L_0]\!] = G_0$.

The transformation is sound in the sense that L_0 and L have the same λ-term as their unfolding.

The procedure for deciding whether two given $\lambda_{\mathsf{letrec}}$-terms L_1 and L_2 are unfolding-equivalent computes their term graph interpretations $[\![L_1]\!]$ and $[\![L_2]\!]$, and checks whether these term graphs are bisimilar.

We also provide an implementation.

3.2 Preliminaries

Definition 3.2.1 (bisimulation collapse)**.** Let $G = \langle V, \mathrm{lab}, \mathrm{args}, r\rangle$ be a term graph. A *bisimulation collapse* of G is a maximal element in the class $\{G' \mid G \trianglerighteq G'\}$ up to $\sim$, that is, a term graph G'_0 with $G \trianglerighteq G'_0$ such that if $G'_0 \trianglerighteq G''_0$ for some term graph G''_0, then $G''_0 \sim G'_0$. Every two bisimulation collapses of G are isomorphic. This justifies the common abbreviation of saying that 'the bisimulation collapse' of G is unique up to isomorphism.

3.3 Introduction

§ 3.3.1 (sharing by letrec)**.** Explicit sharing in pure functional programming languages is typically expressed by means of the letrec-construct, which facilitates cyclic definitions (see also § 0.1.2). For the programmer the letrec-construct offers the possibility to write a program more compactly by utilising subterm sharing. letrec-expressions bind subterms to variables; these variables then denote occurrences of the respective subterms and can be used anywhere inside of the letrec-expression (also recursively). In this way, instead of repeating a subterm multiple times, a single definition can be given which is then referenced from multiple positions.

Example 3.3.2 (horizontal sharing)**.** Consider the λ-term $(\lambda x.\, x)\ (\lambda x.\, x)$ with two occurrences of the subterm $\lambda x.\, x$. These occurrences can be shared as done in the $\lambda_{\mathsf{letrec}}$-term $\mathsf{let}\ id = \lambda x.\, x\ \mathsf{in}\ id\ id$. Obviously it holds: $[\![\mathsf{let}\ id = \lambda x.\, x\ \mathsf{in}\ id\ id]\!]_{\boldsymbol{\lambda}} = (\lambda x.\, x)\ (\lambda x.\, x)$

As let-expressions permit cyclic definitions, sharing can not only occur horizontally but also vertically.

Example 3.3.3 (vertical sharing)**.** Consider the $\lambda_{\mathsf{letrec}}$-terms L and P and the λ-term M from example 1.2.5:

$$L := \lambda f.\, \mathsf{let}\ r = f\ r\ \mathsf{in}\ r \qquad\qquad M := \lambda f.\, f\ (f\ (\ldots))$$
$$P := \lambda f.\, \mathsf{let}\ r = f\ (f\ r)\ \mathsf{in}\ r$$

Both L and P have M as their infinite unfolding: $[\![L]\!]_{\boldsymbol{\lambda}} = [\![P]\!]_{\boldsymbol{\lambda}} = M$. Note that L represents M in a more compact way than P. It is intuitively clear that there is no $\lambda_{\mathsf{letrec}}$-term that represents M more compactly than L. So L can be called a 'maximally shared form' of P (and of M).

Remark 3.3.4 (twisted sharing)**.** Besides horizontal and vertical sharing also a hybrid form of sharing can occur, a sort of superimposition of both kinds, which is called 'twisted sharing' in [8, Definition 4.1.7].

§ 3.3.5 (dynamic vs. static sharing)**.** In the context of functional programming 'sharing' can refer to two different – albeit related – notions. Here, we call them *static* and *dynamic* sharing, while in the literature about static and dynamic sharing, this distinction is usually not made explicitly.

Static sharing simply refers to a trait of some (or graph) languages in which a term (graph) with multiple occurrences of the same subterm (subgraph) can also be written more compactly, where the subterm (subgraph) is written out only once and referenced from multiple points. Static sharing is possible in most programing languages and most optimising compilers perform common subexpression elimination at compile time to increase sharing. $\boldsymbol{\lambda}_{\mathsf{letrec}}$ and the graph formalisms from the previous chapter with their unfolding semantics are of course typical examples of languages with static sharing. This thesis focuses (almost exclusively) on static sharing. *Dynamic sharing* refers to the degree of 'static' sharing an evaluator is able to maintain during evaluation. This is an important issue because unsharing is an integral part of evaluation. A shared function typically behaves differently in different contexts (i.e. with different input), therefore it is impossible to evaluate the entire function only

once if the result is required for two different inputs. However some portion of the computation may very well be shared. The degree of sharing of an evaluator refers to how clever it performs unsharing, and therefore how much of the computation can be shared. Terms as 'call-by-need', 'full laziness' [53], 'complete laziness', and 'optimal evaluation' [39, 4] all refer to dynamic sharing. An overview can be found in [52, 3.4]. While our maximal sharing is a priori a method for increasing static, not dynamic, sharing, we do envisage applying it as part of an evaluator, collapsing the program's graph representation periodically at run-time (see section 3.10).

§ 3.3.6 ('maximal sharing' in ATERM)**.** The term 'maximal sharing' stems from work on the ATERM library [9]. It describes a technique for minimising memory usage when representing a set of terms in a first-order term rewrite system (TRS). The terms are kept in an aggregate directed acyclic graph by which their syntax trees are shared as much as possible. Thereby terms are created only if they are entirely new; otherwise they are referenced by pointers to roots of sub-dags. Our use of the expression 'maximal sharing' is inspired by that work, but our results generalise that approach in the following ways:

- Instead of first-order terms we consider terms in higher-order languages.
- Since letrec can express cyclic sharing, we interpret terms as cyclic graphs instead of just dags.
- We increase sharing by bisimulation collapse instead of by identifying isomorphic sub-dags.

§ 3.3.7 (common subexpression elimination)**.** ATERM only checks for equality of subexpressions. Therefore it only introduces horizontal sharing (for a definition see [8]) and implements a form of *common subexpression elimination (CSE)* [33, 14.7.2]. Our approach is stronger than CSE: while example 3.3.2 can be handled by CSE, this is not the case for example 3.3.3. In contrast to CSE, our approach increases also vertical and twisted sharing (see remark 3.3.4).

3.4 Overview: Methods and Formalisms

Here we will quickly introduce mathematical symbols for the central formalisms (which are later properly defined) so that we can sketch a complete picture of the algorithms we develop in this chapter.

§ 3.4.1 (graph formalisms)**.** The methods that we introduce in this chapter rely heavily on the term graph formalisms developed in chapter 2. As a main result of chapter 2 we have identified suitable classes of term graphs for representing regular λ-terms. In particular we will use eager-scope λ-ap-ho-term-graphs[1] with variable backlinks and eager-scope λ-term-graphs with variable and scope-delimiter backlinks, however over a slightly modified signature which includes black holes. In this chapter we will use an abbreviated notation for these classes. λ-ap-ho-term-graphs (originally $\mathcal{H}_1^{(\lambda)}$) amended with black holes are now denoted by $\mathcal{H}$. λ-term-graphs (originally $\mathcal{T}_{1,2}^{\lambda}$) amended with black holes are now denoted by $\mathcal{T}$.

§ 3.4.2 (graph semantics)**.** We provide a higher-order graph semantics $[\![\cdot]\!]_{\mathcal{H}}$ for interpreting $\lambda_{\mathsf{letrec}}$-terms as eager-scope λ-ap-ho-term-graphs. Together with the correspondence from theorem 2.7.20 it induces a first-order graph semantics $[\![\cdot]\!]_{\mathcal{T}}$. Specifically we use the mapping $\mathrm{G}_{1,2}$ from λ-ap-ho-term-graphs to λ-term-graphs from proposition 2.7.18, which (amended to handling black holes) we call $\mathcal{HT}$ in this chapter.

§ 3.4.3 (readback)**.** In order to be able to compute the $\lambda_{\mathsf{letrec}}$-term that a λ-term-graph stands for, we provide a readback function *rb* from λ-term-graphs to $\lambda_{\mathsf{letrec}}$-terms with the property that it is a right inverse of $[\![\cdot]\!]_{\mathcal{T}}$ up to graph isomorphism. A *readback* function *rb* from λ-term-graphs to $\lambda_{\mathsf{letrec}}$-terms that, for every λ-term-graph G, computes a $\lambda_{\mathsf{letrec}}$-term L from the set of $\lambda_{\mathsf{letrec}}$-terms that have G as their first-order interpretation via $[\![\cdot]\!]_{\mathcal{H}}$ and $\mathcal{HT}$ (i.e. a $\lambda_{\mathsf{letrec}}$-term for which it holds that $\mathcal{HT}([\![L]\!]_{\mathcal{H}}) = G$).

§ 3.4.4 (methods)**.** Putting the above formalisms together we obtain the followings methods, illustrated in fig. 3.1):

- *Maximal sharing*: for a given $\lambda_{\mathsf{letrec}}$-term, a maximally shared form can be obtained by collapsing its first-order term graph interpretation, and then reading back the collapse: $rb \circ |\!\downarrow \circ [\![\cdot]\!]_{\mathcal{T}}$
- *Unfolding equivalence*: for given $\lambda_{\mathsf{letrec}}$-terms L and P, it can be decided whether $[\![L]\!]_{\boldsymbol{\lambda}} = [\![P]\!]_{\boldsymbol{\lambda}}$ by checking whether their term graph interpretations $[\![L]\!]_{\mathcal{T}}$ and $[\![P]\!]_{\mathcal{T}}$ are bisimilar.

See fig. 3.2 for an illustration of the application of the maximal sharing method to the $\lambda_{\mathsf{letrec}}$-terms L and P from example 3.3.3.

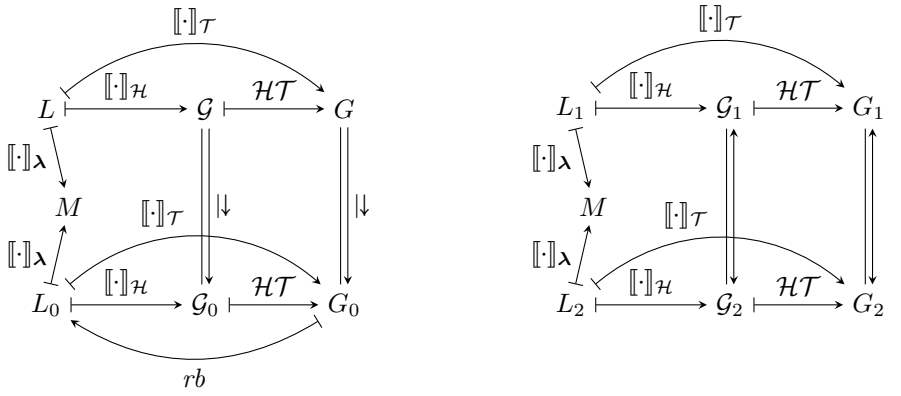

Figure 3.1. On the left: computing the maximally shared form L_0 of a $\lambda_{\mathsf{letrec}}$-term L via bisimulation collapse $|\!\downarrow$. On the right: deciding unfolding equivalence of $\lambda_{\mathsf{letrec}}$-terms L_1 and L_2 via bisimilarity $\leftrightarrows$.

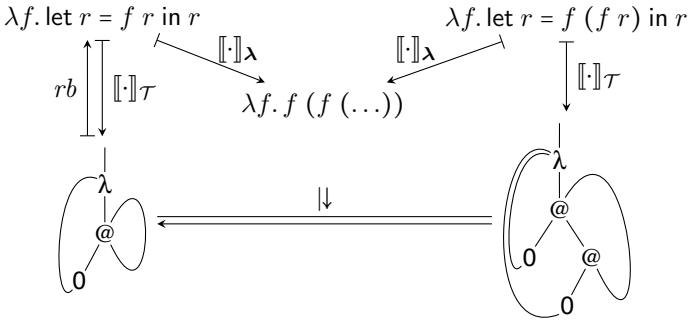

Figure 3.2. Computing the maximally shared version of the term P (on the right) from example 3.3.3 yielding L (on the left).

§ 3.4.5 (correctness and practicality)**.** The correctness of these methods hinges on the fact that the translation and the readback satisfy the following properties:

(P1) $\lambda_{\mathsf{letrec}}$-terms L and P have the same infinite unfolding if and only if the term graphs $[\![L]\!]_{\mathcal{T}}$ and $[\![P]\!]_{\mathcal{T}}$ are bisimilar.

(P2) The class $\mathcal{T}$ of λ-term-graphs is closed under functional bisimulation.

(P3) The readback rb is a right inverse of $[\![\cdot]\!]_{\mathcal{T}}$ up to isomorphism $\sim$, that is, for all term graphs $G \in \mathcal{T}$ it holds: $[\![rb(G)]\!]_{\mathcal{T}} \sim G$.

Furthermore, practicality of these methods depends on the property:

(P4) Translation $[\![\cdot]\!]_{\mathcal{T}}$ and readback rb are efficiently computable.

§ 3.4.6 (applications)**.** Our approach holds promise for a number of practical applications:

- Increasing the efficiency of the execution of programs by transforming them into their maximally shared form at compile-time.
- Increasing the efficiency of the execution of programs by periodically compactifying the program at run time.
- Improving systems for recognising program equivalence.
- Providing feedback to the programmer, along the lines: 'This code has identical fragments and can be written more compactly.'

These and a number of other potential applications are discussed in more detail in section 3.10.

3.5 Interpretion of $\lambda_{\mathsf{letrec}}$-terms as λ-ap-ho-term-graphs

§ 3.5.1 (natural first-order semantics)**.** First we will quickly look at a naive translation of $\lambda_{\mathsf{letrec}}$-terms into first-order term graphs over $\Sigma^{\lambda}_{0_1}$ which we call the *natural first-order semantics* of $\boldsymbol{\lambda}_{\mathsf{letrec}}$ and see that it does not work. We will not give a formal definition of that translation but only the following informal description: consider the syntax tree of a $\lambda_{\mathsf{letrec}}$-term and resolve variable occurrences as $\mathsf{0}$-vertices with variable backlinks; furthermore resolve occurrences of function variables as an edge to the root of the corresponding function's subgraph in the term graph translation.

[1] Note that λ-ho-term-graphs would be just as suitable due to theorem 2.5.11

Example 3.5.2 (incorrectness of the natural first-order semantics)**.** For the terms L, L_1, L_2 below, it holds that $[\![L_1]\!]_{\boldsymbol{\lambda}} = [\![L]\!]_{\boldsymbol{\lambda}} \neq [\![L_2]\!]_{\boldsymbol{\lambda}}$:

$$\begin{aligned} L_1 &= \mathsf{let}\ f = \lambda x.\,(\lambda y.\,f\ y)\ x\ \mathsf{in}\ f \\ L &= \mathsf{let}\ f = \lambda x.\,f\ x\ \mathsf{in}\ f \\ L_2 &= \mathsf{let}\ f = \lambda x.\,(\lambda y.\,f\ x)\ x\ \mathsf{in}\ f \end{aligned}$$

Their term graph interpretation under the natural first-order semantics G_1, G, and G_2 are however all bisimilar:

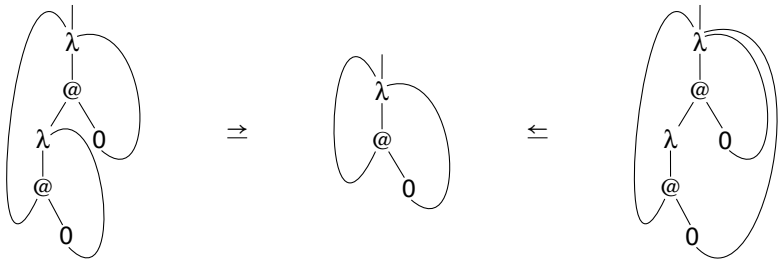

G_1 $\qquad$ G $\qquad$ G_2

This violates property (P1) as under this semantics bisimilarity does not guarantee unfolding equivalence. Therefore the natural first-order semantics is incorrect.

Let us consider the translation of the above example into λ-ap-ho-term-graphs, before we give a formal definition thereof.

Example 3.5.3 (the λ-ap-ho-term-graphs of the terms in example 3.5.2)**.**

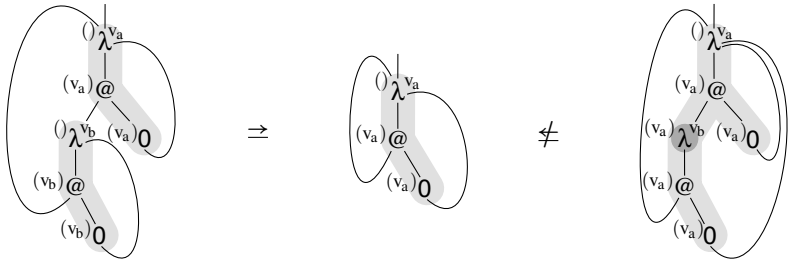

$[\![L_1]\!]_{\mathcal{H}}$ $\qquad$ $[\![L]\!]_{\mathcal{H}}$ $\qquad$ $[\![L_2]\!]_{\mathcal{H}}$

We see that the abstraction prefixes prevent $[\![L]\!]_{\mathcal{H}}$ and $[\![L_2]\!]_{\mathcal{H}}$ to be bisimilar, and that property (P1) is satisfied in this instance.

§ 3.5.4 (cf. bisimulation on the underlying term graphs)**.** The shortcoming of the natural first-order semantics is reflected in proposition 2.6.5 which states that higher-order term graphs are not closed under functional bisimulation on their underlying term graphs.

§ 3.5.5 (λ-ap-ho-term-graphs with black holes)**.** Before we finally define the higher-order term graph semantics $[\![\cdot]\!]_{\mathcal{H}}$ for $\lambda_{\mathsf{letrec}}$-terms, we first have to slightly adapt the graph formalism. The reason is that in chapter 2 we introduced higher-order term graphs as a representation for strongly regular λ-terms, while here we want them to represent $\lambda_{\mathsf{letrec}}$-terms. We know these two classes to coincide (theorem 1.6.28), but only if we exclude meaningless $\lambda_{\mathsf{letrec}}$-terms (see § 0.3.7). Here, however, we want to be able to handle all $\lambda_{\mathsf{letrec}}$-terms. Therefore in this chapter we need to consider an amended definition of λ-ap-ho-term-graphs over an extended signature $\Sigma^{\lambda}_{\bullet}$ which includes black hole vertices.

Definition 3.5.6 (signature for λ-ap-ho-term-graphs with black holes)**.** By $\Sigma^{\lambda}_{\bullet}$ we denote the signature $\Sigma^{\lambda}_{0_1}$ extended by a black-hole symbol $\bullet$ with arity 0, i.e. $\Sigma^{\lambda}_{\bullet} = \{@, \lambda, \mathsf{0}, \bullet\}$ with $\mathrm{ar}(@) = 2$, $\mathrm{ar}(\lambda) = 1$, $\mathrm{ar}(\mathsf{0}) = 1$, and $\mathrm{ar}(\bullet) = 0$.

We also have to amend the correctness conditions definition 2.5.3 for the abstraction-prefix function by a black-hole case.

Definition 3.5.7 (abstraction-prefix function for $\Sigma^{\lambda}_{\bullet}$-term-graphs)**.** As definition 2.5.3 but over signature $\Sigma^{\lambda}_{\bullet}$ and with the following condition added:

$$\mathrm{P}(\bullet) = \epsilon \qquad \text{(black hole)}$$

Definition 3.5.8 (λ-ap-ho-term-graph with black holes)**.** As definition 2.5.4, but over signature $\Sigma^{\lambda}_{\bullet}$ and using definition 3.5.7.

Terminology 3.5.9 (λ-ap-ho-term-graph := λ-ap-ho-term-graph with black holes)**.** Henceforth, if we speak of λ-ap-ho-term-graphs, we refer to this specific variant with black holes and variable backlinks. All the relevant properties that hold for ordinary λ-ap-ho-term-graphs carry over; accounting for the black holes is generally easy.

§ 3.5.10 (translating $\lambda_{\mathsf{letrec}}$-terms into λ-ap-ho-term-graphs)**.** In order to interpret a $\lambda_{\mathsf{letrec}}$-term L as a λ-ap-ho-term-graph, the translation rules $\mathcal{R}$ from fig. 3.3 are applied to a 'translation box' $\boxed{(\star\{\})\ L}$. It contains L furnished with a prefix consisting of a dummy variable $\star$ annotated with an empty set $\{\}$ of function variables. The translation process proceeds by induction on the

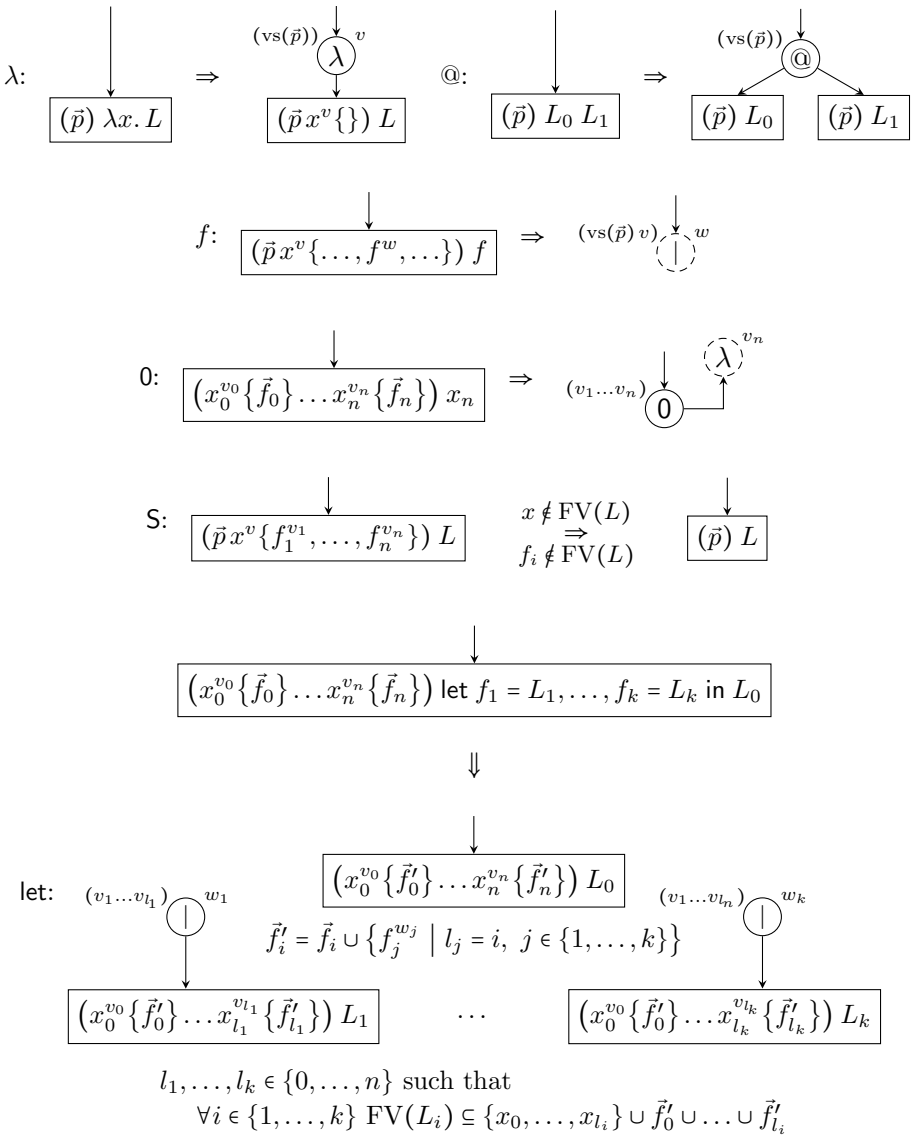

Figure 3.3. Translation rules $\mathcal{R}$ for interpreting $\lambda_{\mathsf{letrec}}$-terms as λ-ap-ho-term-graphs. See section 3.5 for explanations.

syntactical structure $\lambda_{\mathsf{letrec}}$-expression. Ultimately, a term graph G over $\Sigma_{\bullet}^{\lambda}$ is produced, together with a correct abstraction-prefix function for G.

For reading the rules $\mathcal{R}$ in fig. 3.3 correctly, take notice of the explanations below. For illustration of their application, please refer to appendix A where several $\lambda_{\mathsf{letrec}}$-terms are translated into λ-ap-ho-term-graphs.

- A translation box $\boxed{(\vec{p})\ L}$ contains a prefixed, partially decomposed $\lambda_{\mathsf{letrec}}$-term L. The prefix contains a vector $\vec{p}$ of annotated λ-abstractions that have already been translated and whose scope typically extends into L. Every variable in the prefix is annotated with a set of function variables that are defined at its level. There is a special dummy variable $*$ as the very first entry of the prefix that carries function variables for top-level function definitions, i.e. definitions that do not reside under any enclosing λ-abstraction. The λ-rule strips off an abstraction from the body of the expression, and pushes the abstraction variable into the prefix, which initially contains an empty set of function variables.

- Names of abstraction vertices are indicated to the right, and abstraction-prefixes to the left of the created vertices. In the following example the λ-abstraction vertex v has the abstraction-prefix $\vec{p}$:

$$\overset{(\vec{p})}{\textcircled{\lambda}}\ v$$

- In order to refer to the vertices in the prefix we use the following notation: $\mathrm{vs}(\vec{p}) = v_1 \ldots v_n$ given that $\vec{p} = \star\{\vec{f}_0\}\, x_1^{v_1}\{\vec{f}_1\} \ldots x_n^{v_n}\{\vec{f}_n\}$.

- Vertices drawn with dashed lines have been created earlier during the translation, and are referenced by new edges in the current translation step.

- $\mathrm{FV}(L)$ is the set of free variables in L.

- The letrec-rule for translating letrec-expressions creates a box for the body as well as for each of its function definitions. For each function definition an *indirection vertex* is created. These vertices guarantee the well-definedness of the process when it translates meaningless bindings such as $f = f$, or $g = h$, $h = g$, which would otherwise give rise to loops without vertices; the result would not be a term graph. Indirection vertices are eliminated by an erasure process at the end: Every indirection vertex that does not point to itself is removed, redirecting all incoming edges to its successor vertex. Finally every loop on a single indirection vertex is replaced by

$$\frac{(\vec{p}\, x^{v}\{\})\ L}{(\vec{p})\ \lambda x.\, L}\ \lambda \qquad \frac{(\vec{p})\ L_0 \qquad (\vec{p})\ L_1}{(\vec{p})\ L_0\ L_1}\ @$$

$$\frac{}{(\vec{p}\, x^{v}\{\ldots, f^{w}, \ldots\})\ f}\ \mathrm{rec} \qquad \frac{}{\left(x_0^{v_0}\{\vec{f}_0\}\ldots x_n^{v_n}\{\vec{f}_n\}\right)\ x_n}\ \mathsf{0}$$

$$\frac{(\vec{p})\ L}{(\vec{p}\, x^{v}\{f_1^{v_1}, \ldots, f_n^{v_n}\})\ L}\ \mathsf{S}\quad (\text{if } x \notin \mathrm{FV}(L) \text{ and } f_i \notin \mathrm{FV}(L))$$

$$\frac{\left(x_0^{v_0}\{\vec{f}'_0\}\ldots x_{l_0}^{v_{l_0}}\{\vec{f}'_{l_0}\}\right)\ L_0\ \ldots\ \left(x_0^{v_0}\{\vec{f}'_0\}\ldots x_{l_k}^{v_{l_k}}\{\vec{f}'_{l_k}\}\right)\ L_k}{\left(x_0^{v_0}\{\vec{f}_0\}\ldots x_n^{v_n}\{\vec{f}_n\}\right)\ \mathsf{let}\ f_1, \ldots, f_k\ \mathsf{in}\ L_0}\ \mathsf{let}$$

with $l_0 = n$ and $l_1, \ldots, l_k$, $\vec{f}'_0, \ldots, \vec{f}'_n$ as in rule let in fig. 3.3

Figure 3.4. Alternative formulation as inference rules of the translation rules in fig. 3.3 for the interpretation of $\lambda_{\mathsf{letrec}}$-terms as λ-ap-ho-term-graphs.

a *black hole* vertex with an empty abstraction prefix that represents a meaningless binding.

- The let-rule is non-deterministic as there is some freedom on choosing the prefix-lengths used for the translation of each function definition. Say, a function f does not use the rightmost variable x in the current abstraction prefix. Then this freedom allows the translation to either remove x from the prefix within the translation of f's definition, or alternatively at every use site of f outside of f's translation. This freedom is limited by the scoping condition at the bottom of the rule: function definitions may only depend on variables and functions that occur in their respective prefix. In this context also note, that the choice of the prefix-lengths used for some function f also determines the position of f within the prefixes used in the translation of the other functions (and the body of the let-expression).

Definition 3.5.11 ($\mathcal{R}$-generated term graphs)**.** We say that a term graph G over $\Sigma^{\lambda}_{\bullet}$ and an abstraction-prefix function P is $\mathcal{R}$-*generated from* a $\lambda_{\mathsf{letrec}}$-term L if G and P are obtained by applying the rules $\mathcal{R}$ from fig. 3.3 to $\boxed{(\star\{\})\ L}$.

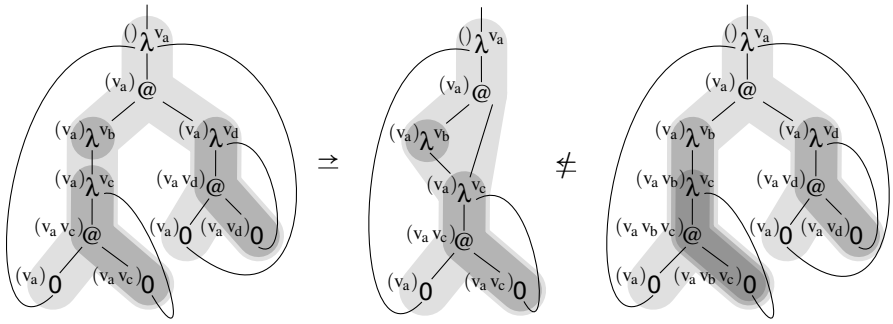

Figure 3.5. Translation of $\lambda a.\,(\lambda b.\,\lambda c.\,a\ c)\ (\lambda d.\,a\ d)$ with eager scope-closure (left), and with lazy scope-closure (right). While on the left four vertices can be shared, on the right only a single variable occurrence can be shared.

§ 3.5.12 (inference rule formulation of $\mathcal{R}$)**.** See also fig. 3.4 for inference rules that correspond to the deconstruction of prefixed terms in $\mathcal{R}$.

Proposition 3.5.13 (λ-ap-ho-term-graph translation of $\lambda_{\mathsf{letrec}}$-terms)**.** Let L be a $\lambda_{\mathsf{letrec}}$-term. Suppose that a term graph G over $\Sigma^{\lambda}_{\bullet}$, and an abstraction-prefix function P are $\mathcal{R}$-generated from L. Then P is a correct abstraction-prefix function for G, and consequently, G and P together form a λ-ap-ho-term-graph.

§ 3.5.14 (non-determinism in $\mathcal{R}$)**.** There are two sources of non-determinism in this translation: The S-rule for shortening prefixes can be applicable at the same time as other rules. And the let-rule does not fix the lengths $l_1, \ldots, l_k$ of the abstraction prefixes used in the translations of the function definitions of the let-expression. Neither kind of non-determinism affects the underlying term graph that is produced, but induces different abstraction-prefix functions, and thus different λ-ap-ho-term-graphs.

Interpretation as eager-scope λ-ap-ho-term-graphs

§ 3.5.15 (eager-scope closure induces a higher degree of sharing)**.** Of the different translations due to § 3.5.14 we are most interested in the one using the shortest possible abstraction prefixes, thus the translation yielding eager-scope λ-ap-ho-term-graphs (definition 2.10.5). The reason for this choice is illustrated in fig. 3.5: eager-scope closure allows for more sharing.

Also, we will call a *translation process* 'eager-scope' if it resolves the non-determinism in $\mathcal{R}$ in such a way that it always yields eager-scope λ-ap-ho-term-graphs. In order to obtain an eager-scope translation we have to consider the following aspects.

§ 3.5.16 (garbage removal)**.** In the presence of garbage – unused function bindings – a translation process cannot be eager-scope. Consider the term $\lambda x.\, \lambda y.\, \mathsf{let}\ f = x\ \mathsf{in}\ y$. The variable x occurs solely in the unused binding $f = x$, which prevents the application of the S-rule, and hence the closure of the scope of λx, directly below λx. Therefore we henceforth assume that *all unused function bindings are removed* prior to applying the rules $\mathcal{R}$.

§ 3.5.17 (short enough prefix lengths in the let-rule)**.** For obtaining an eager-scope translation we will stipulate that the S-rule is applied eagerly, i.e. it is given precedence over the other rules. This is clearly necessary for keeping the abstraction prefixes minimal. But how do we choose the prefix lengths $l_1, \ldots, l_k$ in the let-rule? The prefix lengths l_i determine at which position a binding $f_i = L_i$ is inserted into the abstraction prefixes. Therefore l_i may not be chosen too short; otherwise a function f depending on a function g may end up to the right of g, and hence may be removed from the prefix by the S-rule prematurely, preventing completion of the translation. Yet simply choosing $l_i = n$ may prevent scopes from being minimal. For example, when translating the term $\lambda a.\, \lambda b.\, \mathsf{let}\ f = a\ \mathsf{in}\ a\ a\ (f\ a)\ b$, it is crucial to allow shorter prefixes for the binding than for the body. This is illustrated in fig. 3.6 where the graph on the left does not have eager scope-closure even if the S-rule is applied eagerly. Consequently the opportunity for sharing the lower application vertices is lost.

§ 3.5.18 (required variable analysis)**.** For choosing the prefixes in the let-rule correctly, the translation process must know for each function binding which λ-variable are 'required' on the right-hand side of its definition. For this we use an analysis obtaining the required variables for positions in a $\lambda_{\mathsf{letrec}}$-term as employed by algorithms for λ-lifting [32, 14]. The term 'required variables' was coined by Morazán and Schultz [42]. A λ-variable x is called *required at a position* p in a $\lambda_{\mathsf{letrec}}$-term L if x is bound by an abstraction above p, and has a free occurrence in the complete unfolding of L below p (also function variables from above p are unfolded). The required variables at position p in L can be computed as those λ-variable with free occurrences that are reachable from p by a downwards traversal with the stipulations: on encountering a let-expression the body is entered; when encountering a function variable the

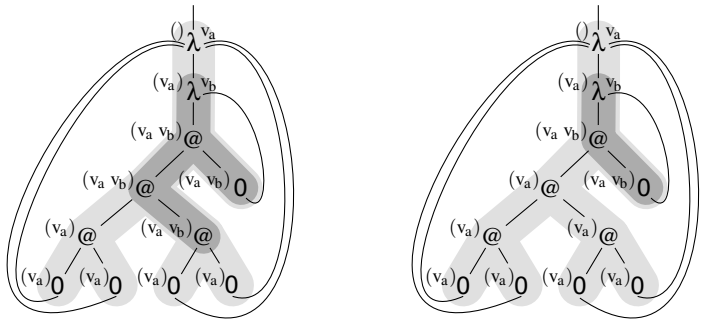

Figure 3.6. Translation of $\lambda a.\, \lambda b.\, \mathsf{let}\ f = a\ \mathsf{in}\ a\ a\ (f\ a)\ b$ with equal (left) and with minimal prefix lengths (right) in the let-rule. See also **??**.

traversal continues at the right-hand side of the corresponding function binding (even if it is defined above p).

With the result of the required variable analysis at hand, we now define properties of the translation process that can guarantee that the resulting λ-ap-ho-term-graph is eager-scope.

Definition 3.5.19 (eager-scope and minimal-prefix generated)**.** Let L be a $\lambda_{\mathsf{letrec}}$-term, and let $\mathcal{G}$ be a λ-ap-ho-term-graph.

We say that $\mathcal{G}$ is *eager-scope* $\mathcal{R}$-generated from L if $\mathcal{G}$ is $\mathcal{R}$-generated from L by a translation process with the following property: for every translation box reached during the process with label $(\lambda\vec{p}\, x^v[f])\ P$, where P is a subterm of L at position q, it holds that if x is not a required variable at q in L, then in the next translation step performed to this box either one of the rules f or let is applied, or the prefix is shortened by the S-rule.

We say that $\mathcal{G}$ is $\mathcal{R}$-generated *with minimal prefixes* from L if $\mathcal{G}$ is $\mathcal{R}$-generated from L by a translation process in which minimal prefix lengths are achieved by giving applications of the S-rule precedence over applications of all other rules, and by always choosing prefixes minimally in applications of the let-rule.

Proposition 3.5.20. Let $\mathcal{G}$ be a λ-ap-ho-term-graph that is $\mathcal{R}$-generated from a garbage-free $\lambda_{\mathsf{letrec}}$-term L. The following statements hold:

(i) If $\mathcal{G}$ is eager-scope $\mathcal{R}$-generated from L, then $\mathcal{G}$ is eager-scope.

(ii) If $\mathcal{G}$ is $\mathcal{R}$-generated with minimal prefixes from L, then $\mathcal{G}$ is eager-scope $\mathcal{R}$-generated from L, hence by (i) $\mathcal{G}$ is eager-scope.

Definition 3.5.21 (higher-order term graph semantics)**.** The semantics $[\![\cdot]\!]_{\mathcal{H}}$ of $\lambda_{\mathsf{letrec}}$-terms as λ-ap-ho-term-graphs is defined as $[\![\cdot]\!]_{\mathcal{H}} : \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}}) \to \mathcal{H}$, $L \mapsto [\![L]\!]_{\mathcal{H}} :=$ the λ-ap-ho-term-graph that is $\mathcal{R}$-generated with minimal prefixes from a garbage-free version of L.

Proposition 3.5.22. For every $\lambda_{\mathsf{letrec}}$-term L, $[\![L]\!]_{\mathcal{H}}$ is eager-scope.

In preparation of proving property (P1) in section 3.6, we establish that the semantics $[\![\cdot]\!]_{\mathcal{H}}$ is correct with respect to the unfolding semantics of $\boldsymbol{\lambda}_{\mathsf{letrec}}$.

Theorem 3.5.23 (correctness of $[\![\cdot]\!]_{\mathcal{H}}$)**.** $[\![L_1]\!]_{\boldsymbol{\lambda}} = [\![L_2]\!]_{\boldsymbol{\lambda}}$ if and only if $[\![L_1]\!]_{\mathcal{H}} \leftrightarrows [\![L_2]\!]_{\mathcal{H}}$, for all $\lambda_{\mathsf{letrec}}$-terms L_1 and L_2.

Sketch of Proof. For the proof we introduce a class of λ-ap-ho-term-graphs that have tree form, i.e. they only contain variable backlinks but no other backlinks. We denote that class by $\mathcal{H}_T \subset \mathcal{H}$. Every $\mathcal{G} \in \mathcal{H}$ has a unique 'tree unfolding' $\mathrm{Tree}(\mathcal{G}) \in \mathcal{H}_T$. We make use of the following statements. For all $L, L_1, L_2 \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$, $M, M_1, M_2 \in \mathrm{Ter}(\boldsymbol{\lambda}_{\bullet})$, $\mathcal{G}, \mathcal{G}_1, \mathcal{G}_2 \in \mathcal{H}$, and $T, T_1, T_2 \in \mathcal{H}_T$ it can be shown that:

$$L_1 \to_{\triangledown} L_2 \;\Rightarrow\; [\![L_1]\!]_{\mathcal{H}} \leftarrows [\![L_2]\!]_{\mathcal{H}} \tag{3.1}$$

$$L \twoheadrightarrow_{\triangledown} M \text{ (hence } [\![L]\!]_{\boldsymbol{\lambda}} = M) \;\Rightarrow\; [\![L]\!]_{\mathcal{H}} \leftarrows [\![M]\!]_{\mathcal{H}} \tag{3.2}$$

$$[\![M]\!]_{\mathcal{H}} \in \mathcal{H}_T \tag{3.3}$$

$$[\![M_1]\!]_{\mathcal{H}} \sim [\![M_2]\!]_{\mathcal{H}} \;\Rightarrow\; M_1 = M_2 \tag{3.4}$$

$$\mathcal{G} \leftarrows \mathrm{Tree}(\mathcal{G}) \tag{3.5}$$

$$T_1 \leftrightarrows T_2 \;\Rightarrow\; T_1 \sim T_2 \tag{3.6}$$

$$\mathcal{G}_1 \leftrightarrows \mathcal{G}_2 \;\Rightarrow\; \mathrm{Tree}(\mathcal{G}_1) \sim \mathrm{Tree}(\mathcal{G}_2) \tag{3.7}$$

We can use (3.1) for proving (3.2), and we can use (3.5) with (3.6) for proving (3.7). Now for proving the theorem, let L_1 and L_2 be arbitrary $\lambda_{\mathsf{letrec}}$-terms.

For "$\Rightarrow$", suppose $[\![L_1]\!]_{\boldsymbol{\lambda}} = [\![L_2]\!]_{\boldsymbol{\lambda}}$. Let M be the infinite unfolding of L_1 and L_2, i.e., $[\![L_1]\!]_{\mathcal{H}} = M = [\![L_2]\!]_{\mathcal{H}}$. Then by (3.2) it follows $[\![L_1]\!]_{\mathcal{H}} \leftarrows [\![M]\!]_{\mathcal{H}} \rightarrows [\![L_2]\!]_{\mathcal{H}}$, and hence $[\![L_1]\!]_{\mathcal{H}} \leftrightarrows [\![L_2]\!]_{\mathcal{H}}$.

For "$\Leftarrow$", suppose $[\![L_1]\!]_{\mathcal{H}} \leftrightarrows [\![L_2]\!]_{\mathcal{H}}$. Then by (3.7) it follows that $\mathrm{Tree}([\![L_1]\!]_{\mathcal{H}}) \sim \mathrm{Tree}([\![L_2]\!]_{\mathcal{H}})$. Let $M_1, M_2 \in \mathrm{Ter}(\boldsymbol{\lambda}_{\bullet})$ be the infinite unfoldings of L_1 and L_2, i.e. $M_1 = [\![L_1]\!]_{\boldsymbol{\lambda}}$, and $M_2 = [\![L_2]\!]_{\boldsymbol{\lambda}}$. Then (3.2) together with

the assumption entails $[\![M_1]\!]_{\mathcal{H}} \leftrightarrows [\![M_2]\!]_{\mathcal{H}}$. Since $[\![M_1]\!]_{\mathcal{H}}, [\![M_2]\!]_{\mathcal{H}} \in \mathcal{H}_T$ by (3.3), it follows by (3.6) that $[\![M_1]\!]_{\mathcal{H}} \sim [\![M_2]\!]_{\mathcal{H}}$. Finally, by using (3.4) we obtain $M_1 = M_2$, and hence $[\![L_1]\!]_{\boldsymbol{\lambda}} = M_1 = M_2 = [\![L_2]\!]_{\boldsymbol{\lambda}}$. □

3.6 Interpretion of $\lambda_{\mathsf{letrec}}$-terms as λ-term-graphs

§ 3.6.1 (overview)**.** We have found a way to model sharing in $\boldsymbol{\lambda}_{\mathsf{letrec}}$ by interpreting $\lambda_{\mathsf{letrec}}$-terms as eager-scope λ-ap-ho-term-graphs. We have also identified eager-scope λ-term-graphs as a class of first-order term graphs that faithfully implements (functional) bisimulation on λ-ap-ho-term-graphs. Since we have a higher-order term graph semantics for interpreting $\lambda_{\mathsf{letrec}}$-terms as λ-ap-ho-term-graphs, as well as translation from λ-ap-ho-term-graphs to λ-term-graphs and back, we seem to be able to turn our attention to the readback function as the last required ingredient of the methods described in § 3.4.4. However, we first need to add black holes to the formalisation of λ-term-graphs as done for λ-ap-ho-term-graphs. And then we also need to deal with a few intricacies of λ-term-graphs with regards to readback.

Analogously to § 3.5.5 we extend the signature $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ and the correctness conditions in definition 2.7.5 by a black-hole case.

Definition 3.6.2 (signature for λ-term-graphs with black holes)**.** By $\Sigma^{\lambda}_{\mathsf{S},\bullet}$ we denote the signature $\Sigma^{\lambda}_{0_1,\mathsf{S}_2}$ extended by a black-hole symbol $\bullet$ with arity 0, thus $\Sigma^{\lambda}_{\mathsf{S},\bullet} = \{@, \lambda, \mathsf{0}, \mathsf{S}, \bullet\}$ with $\mathrm{ar}(@) = 2$, $\mathrm{ar}(\lambda) = \mathrm{ar}(\mathsf{0}) = 1$, $\mathrm{ar}(\mathsf{S}) = 2$, and $\mathrm{ar}(\bullet) = 0$.

Definition 3.6.3 (abstraction-prefix function for $\Sigma^{\lambda}_{\mathsf{S},\bullet}$-term-graphs)**.** As definition 2.7.5 but over signature $\Sigma^{\lambda}_{\mathsf{S},\bullet}$ and with the following condition added:

$$\mathrm{P}(\bullet) = \epsilon \qquad \text{(black hole)}$$

Definition 3.6.4 (λ-term-graph with black holes)**.** As definition 2.6.3, but over signature $\Sigma^{\lambda}_{\mathsf{S},\bullet}$ and using definition 3.6.3.

Terminology 3.6.5 (λ-term-graph := λ-term-graph with black holes)**.** Henceforth, if we speak of λ-term-graphs, we refer to this specific variant with black holes and variable and scope-delimiter vertices with backlinks. All the relevant properties that hold for ordinary λ-term-graphs carry over.

Remark 3.6.6 (relaxed black-hole condition)**.** The condition from definition 3.6.3 could also be relaxed to say $\mathrm{P}(\bullet) = w$ where w is an arbitrary word

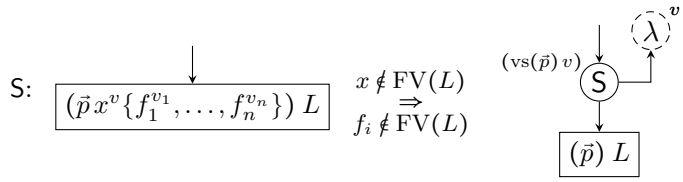

Figure 3.7. Delimiter-vertex producing version of the S-rule from fig. 3.3

of vertices. In that case instead of one single shared function definition of the form $f = f$ one would obtain multiple such definitions, but at most one for each scope.

§ 3.6.7 ($\mathcal{HT}$)**.** $\mathcal{HT}$, the function that maps λ-ap-ho-term-graphs to their corresponding λ-term-graphs is defined as $\mathrm{G}_{1,2}$ from proposition 2.7.18 but amended to also translate black holes (in the obvious way).

§ 3.6.8 (two interpretations as λ-term-graphs)**.** We will consider in fact two interpretations of $\lambda_{\mathsf{letrec}}$-terms as λ-term-graphs: first we define $[\![\cdot]\!]_{\mathcal{T}}^{\min}$ as the composition of $[\![\cdot]\!]_{\mathcal{H}}$ and $\mathcal{HT}$; then we define the semantics $[\![\cdot]\!]_{\mathcal{T}}$ with more fine-grained S-sharing, which is necessary for defining a readback with property (P3).

By composing the interpretation $\mathcal{HT}$ of λ-ap-ho-term-graphs as λ-term-graphs with the λ-ap-ho-term-graph semantics $[\![\cdot]\!]_{\mathcal{H}}$, a semantics of $\lambda_{\mathsf{letrec}}$-terms as λ-term-graphs is obtained. There is, however, a more direct way to define this semantics: by using an adaptation of the translation rules $\mathcal{R}$ in fig. 3.3, on which $[\![\cdot]\!]_{\mathcal{H}}$ is based. For this, let $\mathcal{R}_{\mathsf{S}}$ be the result of replacing the rule S in $\mathcal{R}$ by the version in fig. 3.7. While applications of this variant of the S-rule also shorten the abstraction-prefix, they additionally produce a delimiter vertex.

Here, at the end of the translation process, every loop on an indirection vertex with a prefix of length n is replaced by a chain of n S-vertices followed by a black hole vertex. Note that, while the system $\mathcal{R}_{\mathsf{S}}$ inherits all of the non-determinism of $\mathcal{R}$, the possible degrees of freedom have additional impact on the result, because now they also determine the precise degree of S-vertex sharing. By analogous stipulations as in definition 3.5.19 we define the conditions under which a λ-term-graph is called *eager-scope* $\mathcal{R}_{\mathsf{S}}$-generated, or $\mathcal{R}_{\mathsf{S}}$-generated *with minimal prefixes*, from a $\lambda_{\mathsf{letrec}}$-term. For these notions, statements entirely analogous to proposition 3.5.20 hold.

Definition 3.6.9 ($[\![\cdot]\!]^{\min}_{\mathcal{T}}$)**.** The semantics $[\![\cdot]\!]^{\min}_{\mathcal{T}}$ for $\lambda_{\mathsf{letrec}}$-terms as λ-term-graphs is defined as $[\![\cdot]\!]^{\min}_{\mathcal{T}} : \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}}) \to {}^{\mathrm{eag}}\mathcal{T}$, $L \mapsto [\![L]\!]^{\min}_{\mathcal{T}} :=$ the eager-scope term graph that is $\mathcal{R}_{\mathsf{S}}$-generated with minimal prefixes from a garbage-free version L.

§ 3.6.10 (no S-sharing)**.** For an example, see example 3.6.18 below. In $[\![\cdot]\!]^{\min}_{\mathcal{T}}$, 'min' also indicates that λ-term-graphs obtained via this semantics exhibit minimal (in fact no) sharing (two or more incoming edges) of S-vertices. This is substantiated by the next proposition, in the light of the fact that $\mathcal{HT}$ does not create any shared S-vertices.

Proposition 3.6.11. $[\![\cdot]\!]^{\min}_{\mathcal{T}} = \mathcal{HT} \circ [\![\cdot]\!]_{\mathcal{H}}$.

§ 3.6.12 ($[\![\cdot]\!]^{\min}_{\mathcal{T}}$ is partial)**.** Hence $[\![\cdot]\!]^{\min}_{\mathcal{T}}$ only yields λ-term-graphs without sharing of S-vertices, and therefore its image cannot be all of ${}^{\mathrm{eag}}\mathcal{T}$. As a consequence, we cannot hope to define a readback function rb with respect to $[\![\cdot]\!]^{\min}_{\mathcal{T}}$ that adheres to property (P3), because that requires that the image of the semantics is ${}^{\mathrm{eag}}\mathcal{T}$ in its entirety.

§ 3.6.13 (eager-scope $\mathcal{R}$-generated with maximal prefixes)**.** Therefore we provide an alternative λ-term-graph semantics $[\![\cdot]\!]_{\mathcal{T}}$ with $\mathrm{im}([\![\cdot]\!]_{\mathcal{T}}) = {}^{\mathrm{eag}}\mathcal{T}$. We achieve this by letting the let-binding-structure of the $\lambda_{\mathsf{letrec}}$-term influence the degree of S-sharing as much as possible, while still remaining eager-scope.

By 'letting the let-binding-structure determine the degree of S-sharing' we mean to eliminate the freedom of choice with respect to closing scopes: one can either place a scope-delimiter vertex at the top of the translation of a function binding in which case it is shared, or right above every use site of that function in which case it is not. See fig. 3.8 for illustration.

Definition 3.6.14 (eager-scope $\mathcal{R}$-generated with maximal prefixes)**.** We say that a λ-ap-ho-term-graph $\mathcal{G}$ is *eager-scope $\mathcal{R}$-generated with maximal prefixes* from a $\lambda_{\mathsf{letrec}}$-term L if $\mathcal{G}$ is $\mathcal{R}$-generated from L by a translation process in which in applications of the let-rule the prefixes are chosen maximally, but so that the eager-scope property of the process is not compromised. It can be shown that this condition fixes the prefix lengths per application of the let-rule.

Definition 3.6.15 ($[\![\cdot]\!]_{\mathcal{T}}$)**.** The semantics $[\![\cdot]\!]_{\mathcal{T}}$ for $\lambda_{\mathsf{letrec}}$-terms as λ-term-graphs is defined as $[\![\cdot]\!]_{\mathcal{T}} : \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}}) \to {}^{\mathrm{eag}}\mathcal{T}$, $L \mapsto [\![L]\!]_{\mathcal{T}} :=$ the λ-term-graph that is eager-scope $\mathcal{R}_{\mathsf{S}}$-generated with maximal prefixes from a garbage-free version of L.

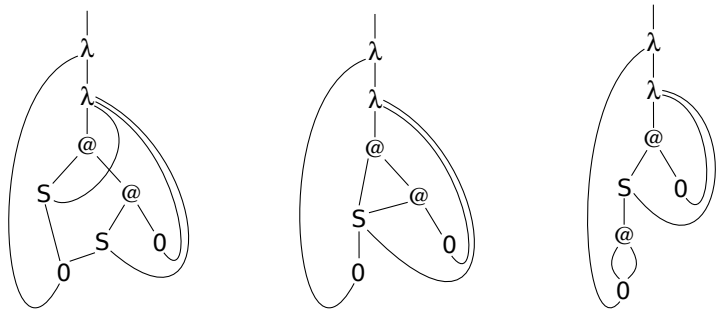

$\lambda x.\,\mathsf{let}\ f = x\ \mathsf{in}\ \lambda y.\, f\ (f\ y)$ $\lambda x.\,\lambda y.\,\mathsf{let}\ f = x\ \mathsf{in}\ f\ (f\ y)$ $\lambda x.\,\lambda y.\,\mathsf{let}\ f = x\ \mathsf{in}\ f\ f\ y$

Figure 3.8. In $[\![\cdot]\!]_{\mathcal{T}}$ the degree of S-sharing is determined by the let-structure as much as possible. Consider the above $\lambda_{\mathsf{letrec}}$-terms and their term graph interpretation w.r.t. $[\![\cdot]\!]_{\mathcal{T}}$.

left : f is bound outside of y's scope. Therefore it is not f's responsibility to close y's scope but the responsibility of f's two use sites.

middle : f is bound within y's scope so it is f's responsibility to close it.

right : f is bound within y's scope. Therefore it should be f's responsibility to close y's scope, but putting the scope delimiter underneath the lower application would violate eager scope-closure.

Proposition 3.6.16. $[\![L]\!]^{\min}_{\mathcal{T}} \trianglerighteq^{S} [\![L]\!]_{\mathcal{T}}$ holds for all $\lambda_{\mathsf{letrec}}$-terms L.

Now due to this, and due to theorem 2.7.20 (iii), the statement of theorem 3.5.23 can be transferred to $\mathcal{T}$, yielding property (P1) for $[\![\cdot]\!]_{\mathcal{T}}$.

Theorem 3.6.17. For all $\lambda_{\mathsf{letrec}}$-terms L_1 and L_2 the following holds: $[\![L_1]\!]_{\boldsymbol{\lambda}} = [\![L_2]\!]_{\boldsymbol{\lambda}}$ if and only if $[\![L_1]\!]_{\mathcal{T}} \leftrightarrows [\![L_2]\!]_{\mathcal{T}}$.

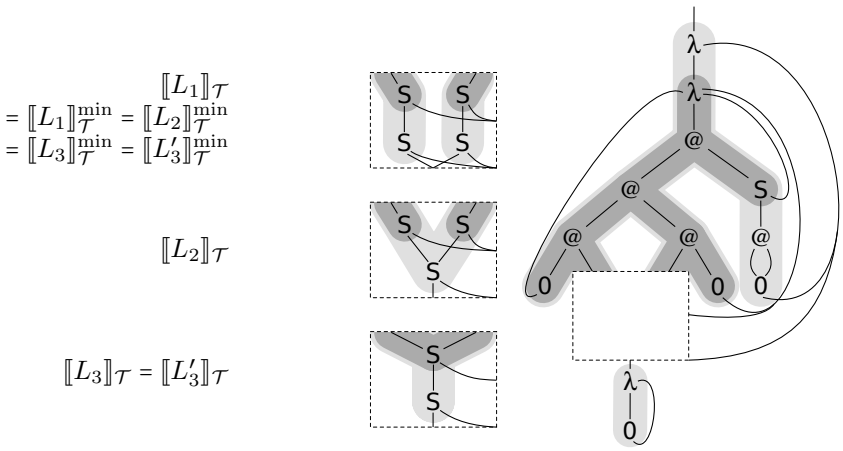

Figure 3.9. Translation of the $\lambda_{\mathsf{letrec}}$-terms from example 3.6.18 with the semantics $[\![\cdot]\!]^{\min}_{\mathcal{T}}$ and $[\![\cdot]\!]_{\mathcal{T}}$. For legibility some backlinks are merged.

Example 3.6.18. Consider the following four $\lambda_{\mathsf{letrec}}$-terms:

$$L_1 = \mathsf{let}\ I = \lambda z.\, z\ \mathsf{in}\ \lambda x.\, \lambda y.\, \mathsf{let}\ f = x\ \mathsf{in}\ ((y\ I)\ (I\ y))\ (f\ f)$$
$$L_2 = \lambda x.\, \mathsf{let}\ I = \lambda z.\, z\ \mathsf{in}\ \lambda y.\, \mathsf{let}\ f = x\ \mathsf{in}\ ((y\ I)\ (I\ y))\ (f\ f)$$
$$L_3 = \lambda x.\, \lambda y.\, \mathsf{let}\ I = \lambda z.\, z,\ f = x\ \mathsf{in}\ ((y\ I)\ (I\ y))\ (f\ f)$$
$$L'_3 = \lambda x.\, \mathsf{let}\ I = \lambda z.\, z\ \mathsf{in}\ \lambda y.\, \mathsf{let}\ f = x,\ g = I\ \mathsf{in}\ ((y\ g)\ (g\ y))\ (f\ f)$$

The three possible fillings of the dashed area in fig. 3.9 depict the translations $[\![L_1]\!]_{\mathcal{T}}$, $[\![L_2]\!]_{\mathcal{T}}$, and $[\![L_3]\!]_{\mathcal{T}} = [\![L'_3]\!]_{\mathcal{T}}$. The translations of the four terms with $[\![\cdot]\!]^{\min}_{\mathcal{T}}$ are identical:
$[\![L_1]\!]^{\min}_{\mathcal{T}} = [\![L_2]\!]^{\min}_{\mathcal{T}} = [\![L_3]\!]^{\min}_{\mathcal{T}} = [\![L'_3]\!]^{\min}_{\mathcal{T}} = [\![L_1]\!]_{\mathcal{T}}$.

See § A.19 for a fully-worked out stepwise translation of L_2 and also **??**.

3.7 Readback of λ-term-graphs

§ 3.7.1 (overview)**.** In this section we describe how from a given λ-term-graph G a $\lambda_{\mathsf{letrec}}$-term L that represents G (i.e. for which $[\![L]\!]_{\mathcal{T}} = G$ holds) can be 'read back'. For this purpose we define a process based on synthesis rules. It defines a readback function from λ-term-graphs to $\lambda_{\mathsf{letrec}}$-terms. We illustrate

this process by an example, formulate its most important properties, and sketch the proof of property (P3).

§ 3.7.2 ($[\![\cdot]\!]_{\mathcal{T}}$ is not invertible)**.** Note that we state as the desired property (P3) of the readback rb to be a right-inverse and not a 'full' inverse of $[\![\cdot]\!]_{\mathcal{T}}$. This is because $[\![\cdot]\!]_{\mathcal{T}}$ is not injective. Consider the $\lambda_{\mathsf{letrec}}$-terms $\lambda x.\,\mathsf{let}\ f = x\ \mathsf{in}\ \lambda y.\,f$ and $\lambda x.\,\lambda y.\,\mathsf{let}\ f = x\ \mathsf{in}\ f$. They only differ in the position of the function definition; this information is lost in the translation; $[\![\cdot]\!]_{\mathcal{T}}$ maps them to the same term graph. In particular, this kind of relocation of function-bindings (called *let-floating* [17]) preserves the λ-term-graph interpretation.

§ 3.7.3 (approach)**.** The idea underlying the definition of the readback procedure is the following: For a given λ-term-graph G, a spanning tree T for G (augmented with a dedicated root node) is constructed that severs cycles due to recursive bindings and variable and scope-delimiter backlinks. Now the spanning tree T facilitates an inductive bottom-up (from the leafs upwards) synthesis process along T, which labels the edges of G with prefixed $\lambda_{\mathsf{letrec}}$-terms. For this process we use local rules (see fig. 3.11) that synthesise labels for incoming edges of a vertex from the labels of its outgoing edges. Eventually the readback of G is obtained as the label for the edge that singles out the root of term graph.

§ 3.7.4 (placement of function definitions)**.** In the design of the readback rules, there is some freedom in where to place the function bindings in the synthesised term (§ 3.7.2). Here, function bindings will be put into a let-expression placed as high up in the term as possible: a binding arising from the term synthesised for a shared vertex v is placed in a let-expression that is created at the enclosing λ-abstraction of v (the rightmost vertex in the abstraction-prefix $\mathrm{P}(v)$ of v).

Definition 3.7.5 (readback of λ-term-graphs)**.** Let $G \in {}^{\mathrm{eag}}\mathcal{T}$ be an eager-scope λ-term-graph. The process of computing the readback of G (a $\lambda_{\mathsf{letrec}}$-term) consists of the following five steps, starting on G:

(Rb-1) Determine the abstraction-prefix function P for G by performing a traversal over G, and associate with every vertex v of G its abstraction-prefix $\mathrm{P}(v)$.

(Rb-2) Add a new vertex on top with label $\top$, arity 1, and empty abstraction prefix. Let G' be the resulting term graph, and P' its abstraction-prefix function.

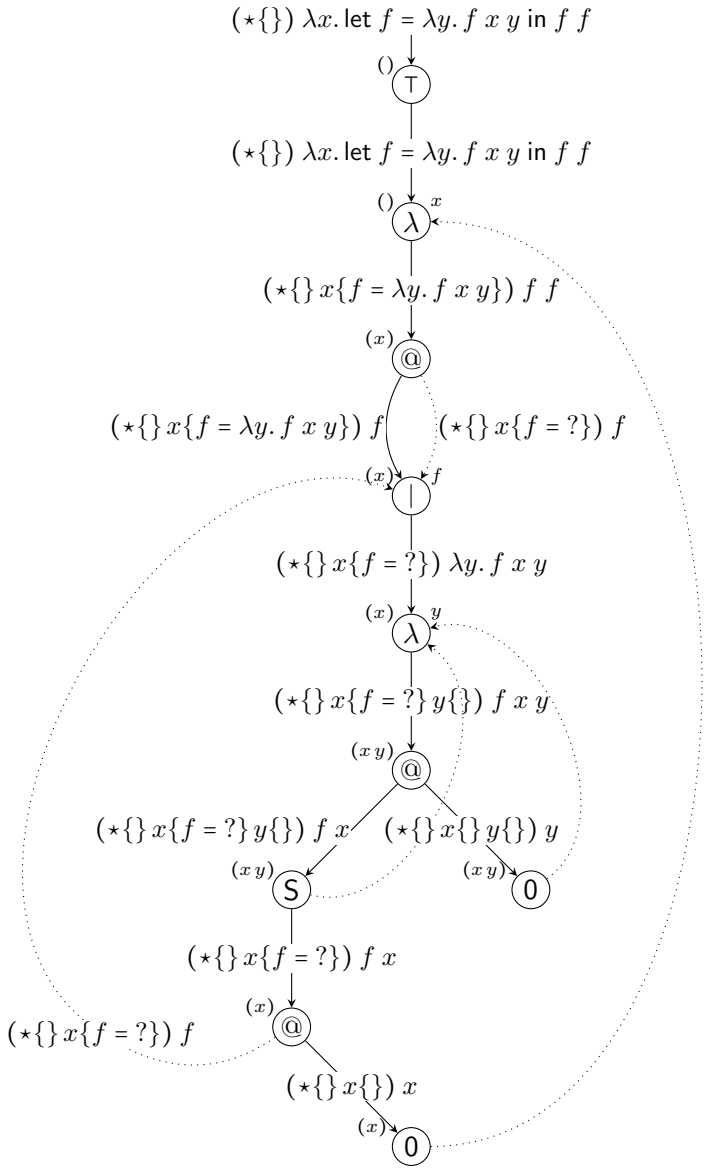

Figure 3.10. Example of the readback synthesis from a λ-term-graph.

(Rb-3) Introduce indirection vertices to organise sharing: For every vertex v of G' with two or more incoming non-variable-backlink edges, add an indirection vertex v_0, redirect the incoming edges of v that are not variable backlinks to v_0, and direct the outgoing edge from v_0 to v. In the resulting term graph G'' only indirection vertices are shared ; their names will be used. Extend P' to an abstraction-prefix function P'' for G'' so that every indirection vertex v_0 gets the prefix of its successor v.

(Rb-4) Construct a spanning tree T'' of G'' by using a depth-first search (DFS) on G''. Note that all variable backlinks and S-backlinks, and some of the recursive backlinks, of G'', are not contained in T'', because they are back-edges of the DFS.

(Rb-5) Apply the readback synthesis rules from fig. 3.11 to G'' with respect to T''. By this a complete labelling of the edges of G'' by prefixed $\lambda_{\mathsf{letrec}}$-terms is constructed. The rules define how the labelling for an incoming edge (on top) of a vertex v is synthesised under the assumption of an already determined labelling of an outgoing edge of (and below) v. If the outgoing edge in the rule does not carry a label, then the labelling of the incoming edge can happen regardless. Note that in these rules:

- full (dotted) lines indicate spanning tree (non-spanning tree) edges; broken lines match either kind;
- abstraction prefixes of vertices are crucial for the $\mathsf{0}$-vertex, and the second indirection vertex rule, where the prefixes in the synthesised terms are created; in the other rules the prefix of the assumed term is used; for indicating a correspondence between a term's and a vertex's abstraction prefix we denote by $\mathrm{vs}(\vec{p})$ the word of vertices occurring in a term's prefix $\vec{p}$;
- the rule for indirection vertices with incoming non-spanning tree edge introduces an unfinished function binding $f = ?$ for f; unfinished bindings are to be filled in later;
- the @-vertex rule applies only if $\mathrm{vs}(\vec{p}_0) = \mathrm{vs}(\vec{p}_1)$; the operation $\vec{\cup}$ used in the synthesised term's prefix builds the union per prefix variable of the pertaining bindings; if the prefixed terms $(\vec{p}_0)\ L_0$ and $(\vec{p}_1)\ L_1$ assumed in this rule contain a yet unfinished function binding $f = ?$ and a completed binding $f = P$ at a λ-variable z, the synthesised term contains the completed binding $f = P$ at z;

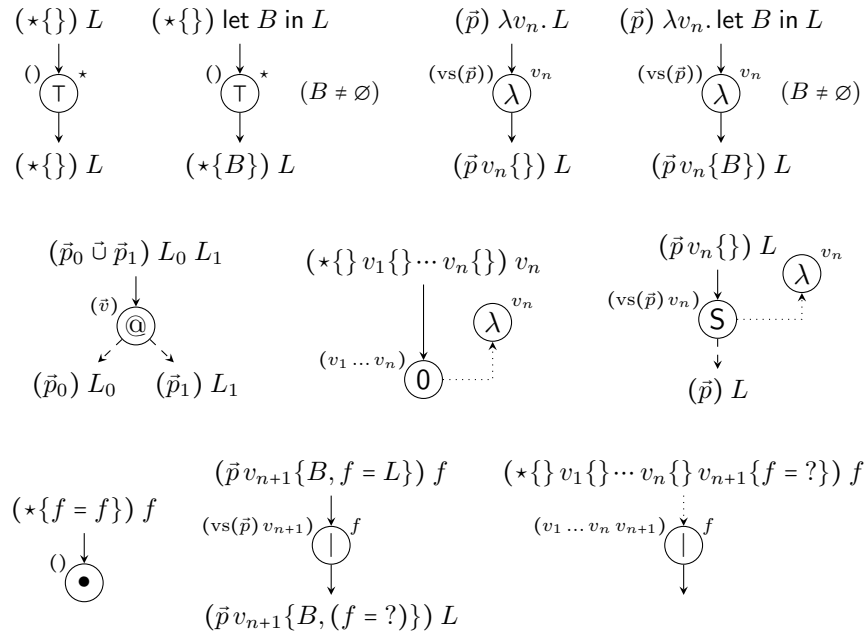

Figure 3.11. Readback synthesis rules for computing a $\lambda_{\mathsf{letrec}}$-term from a λ-term-graph. For explanations, see definition 3.7.5 (Rb-5).

If this process yields the label $(\star\{\})\ L$ at the root of G'', we call L the *readback of* G.

§ 3.7.6 (readback is deterministic)**.** For every edge e the synthesis rule to be applied is uniquely determined by the label of the target vertex v of e together with side conditions for λ, $\top$, and $|$. In the last case it depends on whether e is a spanning-tree edge or not.

Proposition 3.7.7 (readback function)**.** For every λ-term-graph G the process from definition 3.7.5 produces a complete edge labelling of the (modified) term graph, with label $(\star\{\})\ L$ for the root edge, where L is a $\lambda_{\mathsf{letrec}}$-term. Hence it yields L as the readback of G. Thus definition 3.7.5 defines a function $rb : \mathcal{T} \to \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$, the *readback function*.

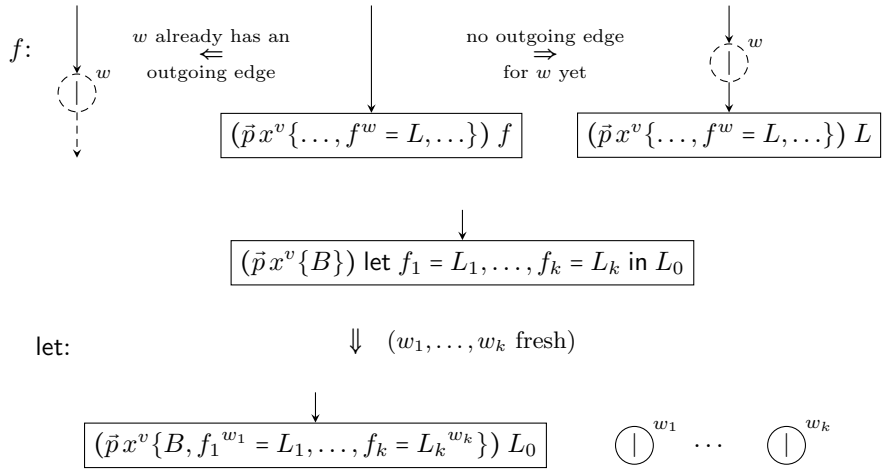

Figure 3.12. Augmented version of two of the translation rules from fig. 3.3 for an alternative definition of the λ-term-graph semantics of $\lambda_{\mathsf{letrec}}$-terms.

Example 3.7.8. See fig. 3.10 for the illustration of the synthesis of the readback from an example λ-term-graph. Full lines depict spanning tree edges, dotted lines depict non-spanning-tree edges.

The following theorem validates property (P3).

Theorem 3.7.9. For all $G \in {}^{\mathrm{eag}}\mathcal{T}$ it holds: $((\llbracket\cdot\rrbracket_{\mathcal{T}} \circ rb))(G) = \llbracket rb(G)\rrbracket_{\mathcal{T}} \sim G$, i.e., rb is a right-inverse of $\llbracket\cdot\rrbracket_{\mathcal{T}}$, and $\llbracket\cdot\rrbracket_{\mathcal{T}}$ a left-inverse of rb, up to $\sim$. Hence rb is injective, and $\llbracket\cdot\rrbracket_{\mathcal{T}}$ is surjective, thus $\mathrm{im}(\llbracket\cdot\rrbracket_{\mathcal{T}}) = {}^{\mathrm{eag}}\mathcal{T}$.

Proof Idea. Graph translation steps can be linked with corresponding readback steps in order to establish that the former roughly reverse the latter. Roughly, because e.g. reversing a λ-readback step necessitates both a λ- and a let-translation step. Therefore, this correspondence holds only for a modification of the translation rules $\mathcal{R}_{\mathsf{S}}$ from fig. 3.3, shown in fig. 3.7 where the rules let (for let-expressions) and f (for occurrences of function variables) are replaced by the locally-operating versions in fig. 3.12. Moreover, the translation rules need to store entire function bindings in the prefix (not just function variables) as the readback rules do. To the augmented rules we also add a initiating rule

$$\frac{}{(\vec{p}\,x^v\{\ldots, f^w = L, \ldots\})\ f}\ \text{rec}$$

$$\frac{(\vec{p}\,x^v\{\ldots, f^w = L, \ldots\})\ L}{(\vec{p}\,x^v\{\ldots, f = L, \ldots\})\ f}\ \text{rec}\quad (w \text{ is fresh})$$

$$\frac{(\vec{p}\,x^v\{B, f_1 = L_1, \ldots, f_k = L_k\})\ L_0}{(\vec{p}\,x^v\{B\})\ \mathsf{let}\ f_1 = L_1, \ldots, f_k = L_k\ \mathsf{in}\ L_0}\ \mathsf{let}$$

Figure 3.13. Formulation of the local translation rules in fig. 3.12 in the form of inference rules.

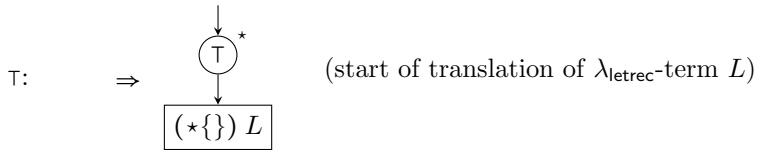

for creating a top vertex. Now the translation of a let-expression does no longer directly spawn translations of the bindings, but the bindings will only be translated later once their calls have been reached during the translation process of the body, or of the definitions of other already translated bindings. Note that in the let-rule in fig. 3.12 function bindings are associated with the rightmost variable in the prefix, which corresponds to choosing $l_i = n$ in the let-rule in fig. 3.3. While such a stipulation does not guarantee the eager-scope translation of every term, it actually does so for all $\lambda_{\mathsf{letrec}}$-terms that are obtained by the readback (on these terms the such defined translation coincides with $[\![\cdot]\!]_{\mathcal{T}}$ from definition 3.5.21).

Please find in fig. 3.14 a graphical argument for the stepwise reversal of readback steps through (augmented) translation steps. This establishes that graph translation steps reverse readback steps, and is the crucial step in the proof of the theorem. The proof uses induction on access paths, and an invariant that relates the eager-scope property localised for a vertex v with the applicability of the S-rule to the readback term synthesised at v. Note that in most cases the sets of function bindings on the left-hand side ($\bar{B}_i$) and the right-hand side (B_i) differ, due to the freedoms of the function definition's positions in the translated $\lambda_{\mathsf{letrec}}$-term (see § 3.7.2).

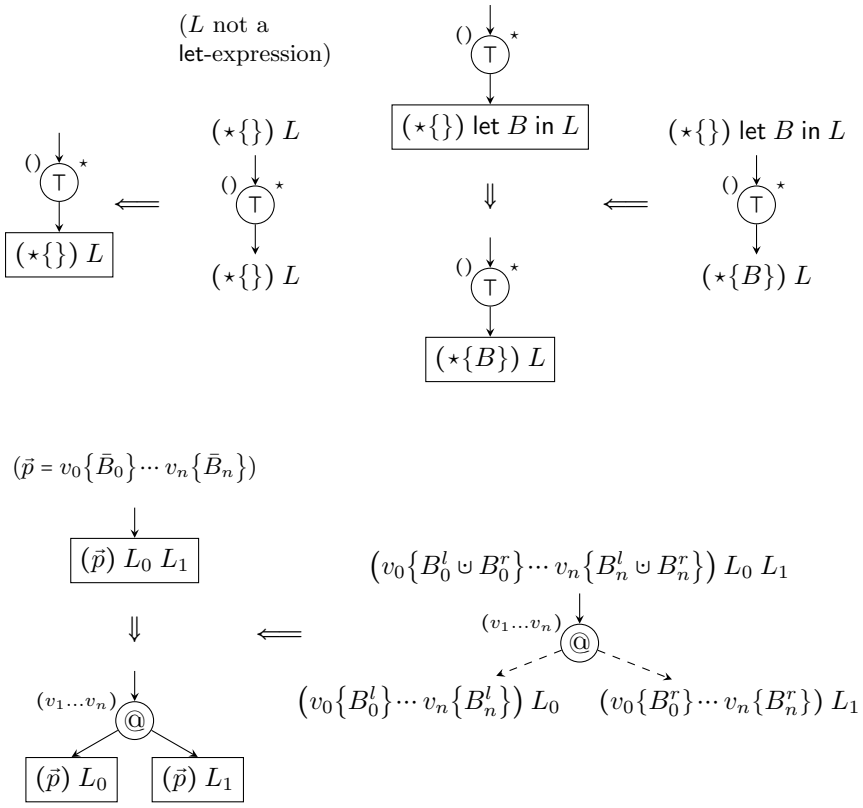

Figure 3.14. Reversal of readback steps through translation steps

Note that $\uplus$ denotes a join of two sets of function definitions where a defined function always overrules an undefined function, such that for instance the following holds: $\{f = x, g = ?\} \uplus \{f = ?, g = ?\} = \{f = x, g = ?\}$. □

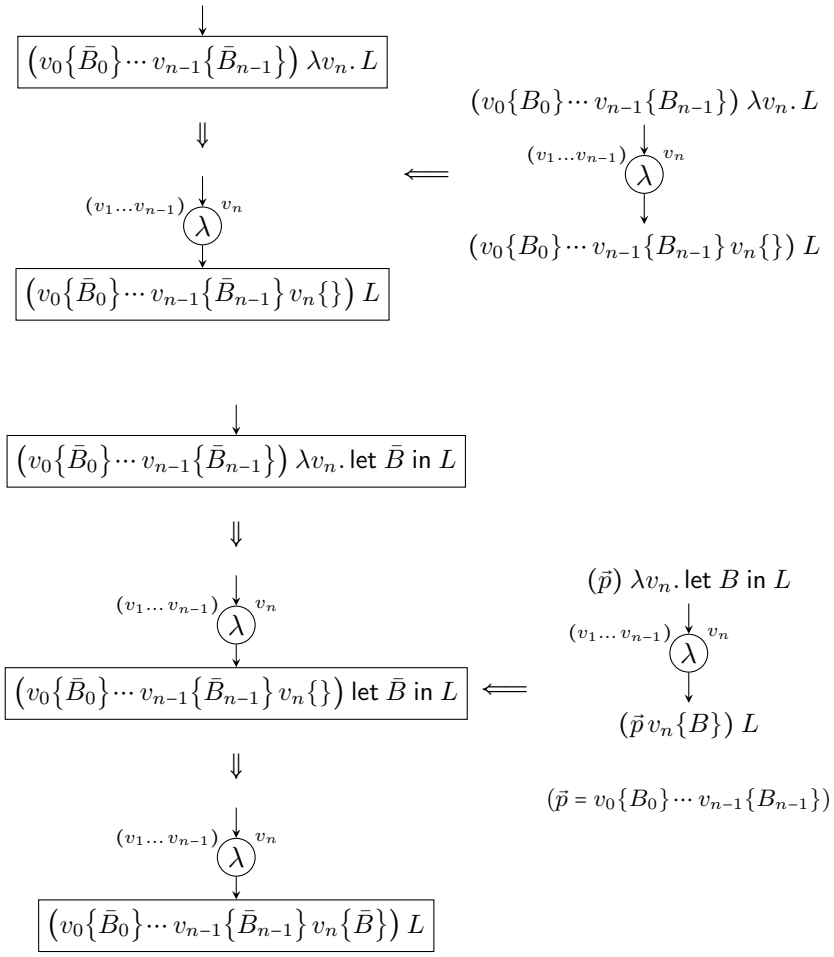

Figure 3.14. (continued) Reversal of readback steps through translation steps

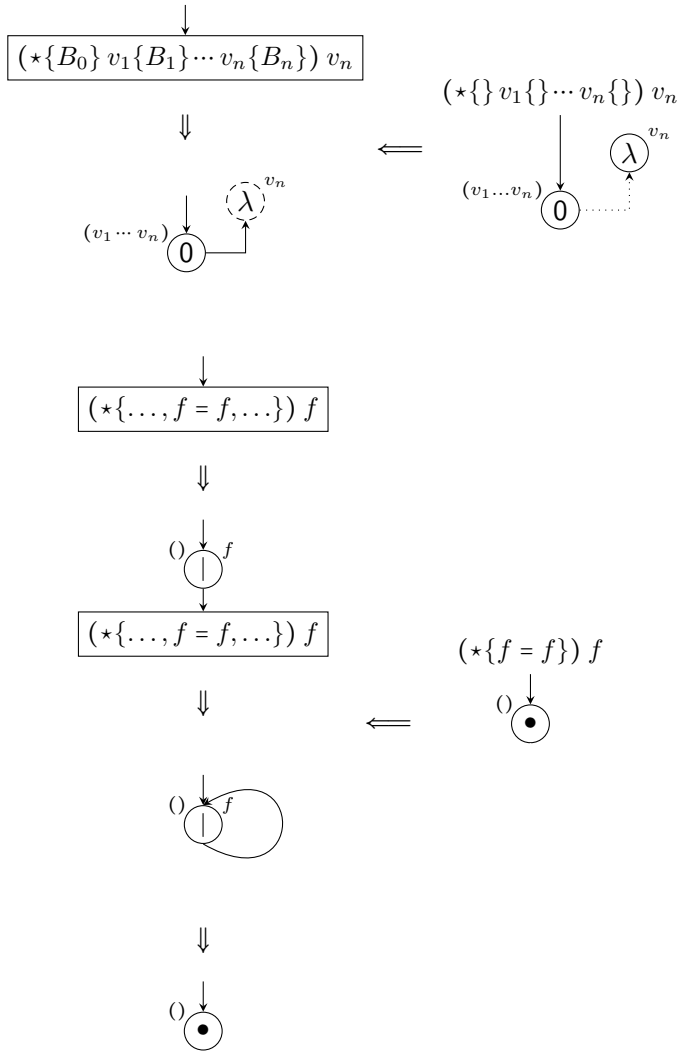

Figure 3.14. (continued) Reversal of readback steps through translation steps

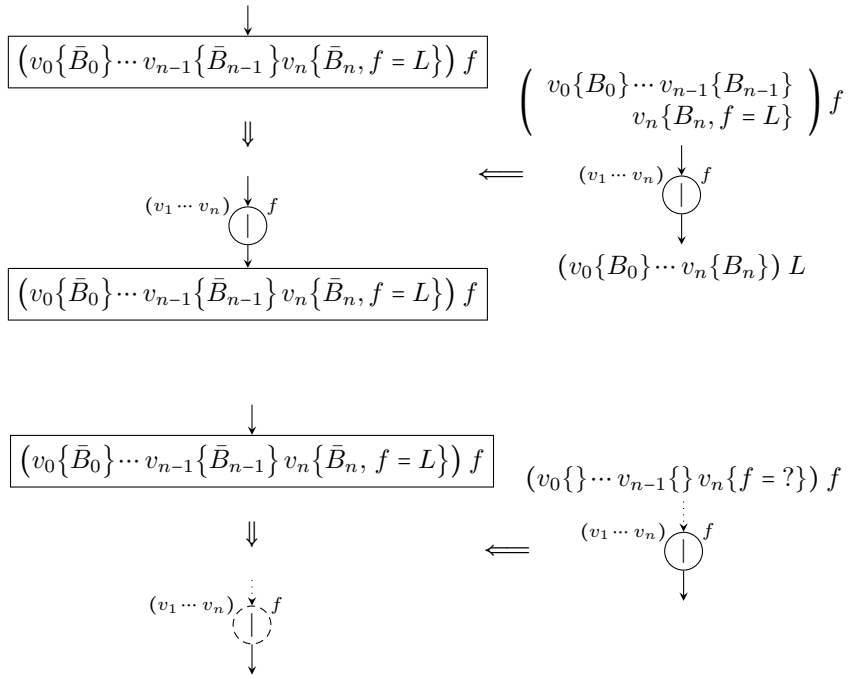

Figure 3.14. (continued) Reversal of readback steps through translation steps

3.8 Complexity analysis

Here we report on the complexity for the individual operations from the previous sections used form implementing maximal sharing and determining unfolding-equivalence.

In the lemma below, (ii) and (v) justify property (P4) of our methods. Items (iii) and (iv) detail the complexity of standard methods when used for computing bisimulation collapse and bisimilarity of λ-term-graphs. Note that first-order term graphs can be modelled by deterministic process graphs, and hence by DFAs [28]. Therefore bisimilarity of term graphs can be computed via language equivalence of the corresponding DFAs (in time $\mathrm{O}(n\alpha(n))$ [43], where α is the quasi-constant *inverse Ackermann function*) and bisimulation collapse via state minimisation of DFAs (in time $\mathrm{O}(n \log n)$) [27].

Consider the finite λ-terms M_n with n occurrences of bindings λx_2:

$$\lambda x_0.\, \lambda x_1.\, x_0\ x_1\ \lambda x_2.\, x_0\ x_1\ \lambda x_1.\, x_0\ x_2\ \lambda x_2.\, x_0\ x_1 \ldots \lambda x_2.\, x_0\ x_1\ x_2$$

$|M_n| \in \mathrm{O}(n)$. But both the transformation of M_n into de-Bruijn notation

$$\lambda.\, \lambda.\, \mathsf{S}(\mathsf{0})\ \mathsf{0}\ \lambda.\, \mathsf{S}^2(\mathsf{0})\ \mathsf{S}(\mathsf{0})\ \lambda.\, \mathsf{S}^3(\mathsf{0})\ \mathsf{S}(\mathsf{0})\ \lambda.\, \mathsf{S}^4(\mathsf{0})\ \mathsf{S}(\mathsf{0}) \ldots \lambda.\, \mathsf{S}^{2n}(\mathsf{0})\ \mathsf{S}(\mathsf{0})\ \mathsf{0}$$

and the rendering of M_n with respect to the eager scope-delimiting strategy:

$$\lambda.\, \lambda.\, \mathsf{S}(\mathsf{0})\ \mathsf{0}\ \lambda.\, \mathsf{S}(\mathsf{S}(\mathsf{0})\ \mathsf{0})\ \lambda.\, \mathsf{S}(\mathsf{S}^2(\mathsf{0})\ \mathsf{0})\ \lambda.\, \mathsf{S}(\mathsf{S}^3(\mathsf{0})\ \mathsf{0}) \ldots \lambda.\, \mathsf{S}(\mathsf{S}^{2n-1}(\mathsf{0})\ \mathsf{0})\ \mathsf{0}$$

have size $\mathrm{O}\big(n^2\big)$.

Figure 3.15. Example of a sequence $\{M_n\}_n$ of finite λ-terms M_n whose translation into λ-term-graphs grows quadratically in the size of M_n.

By $|L|$ we denote the size (number of symbols) of a $\lambda_{\mathsf{letrec}}$-term L. Also we denote by $|G|$ the size (number of vertices) of a term graph G.

Lemma 3.8.1. (i) $|[\![L]\!]_{\mathcal{T}}| \in \mathrm{O}\Big(|L|^2\Big)$ for $L \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$.

(ii) Translating $L \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$ into $[\![L]\!]_{\mathcal{T}} \in \mathcal{T}$ takes time $\mathrm{O}\Big(|L|^2\Big)$.

(iii) Collapsing $G \in \mathcal{T}$ to $G{\downarrow}$ is in $\mathrm{O}(|G| \log |G|)$.

(iv) Deciding bisimilarity of $G_1, G_2 \in \mathcal{T}$ requires time $\mathrm{O}(n\alpha(n))$ for $n = \max\{|G_1|, |G_2|\}$.

(v) Computing the readback $rb(G)$ for a given $G \in \mathcal{T}$ requires time $\mathrm{O}(n \log n)$, for $n = |G|$.

See fig. 3.15 for an example that the size bound in item (i) of the lemma is tight.

Proposition 3.8.2. $|[\![L]\!]_{\mathcal{T}}| \in \mathrm{O}\Big(|L|^2\Big)$ for $\lambda_{\mathsf{letrec}}$-terms L.

Based on this lemma, and on further considerations, we obtain the following complexity statements for our methods.

Theorem 3.8.3. (i) The computation, for a $\lambda_{\mathsf{letrec}}$-term L with $|L| = n$, of a maximally compactified form $(rb \circ |\!\downarrow \circ [\![\cdot]\!]_{\mathcal{T}})(L)$ requires time $\mathrm{O}\big(n^2 \log n\big)$. By using an S-unsharing operation $\mathsf{unsh}_{\mathsf{S}}$ after the collapse, a (typically smaller) $\lambda_{\mathsf{letrec}}$-term $((rb \circ \mathsf{unsh}_{\mathsf{S}} \circ |\!\downarrow \circ [\![\cdot]\!]_{\mathcal{T}}))L$ of size $\mathrm{O}(n \log n)$ can be obtained, with the same time complexity.

(ii) The decision of whether two $\lambda_{\mathsf{letrec}}$-terms L_1 and L_2 are unfolding equivalent requires time $\mathrm{O}\big(n^2\big)\alpha(n)$ for $n = \max\{|L_1|, |L_2|\}$.

3.9 Implementation

We have implemented our methods in Haskell using the Utrecht University Attribute Grammar System. The implementation is available at `http://hackage.haskell.org/package/maxsharing/`. Output produced for a number of examples from this thesis can be found in appendix B. The output includes translations into λ-term-graphs (in DFA-form) according to different semantics, complete derivations w.r.t. these semantics, the bisimulation collapse of the λ-term-graphs and the readback thereof.

3.10 Modifications, extensions and applications

We conclude by describing straightforward modifications, extensions, and promising areas of application for our methods.

Modifications

§ 3.10.1 (implicit sharing of λ-variable)**.** Our method introduces explicit sharing via let for multiple occurences of the same λ-variable. For instance the term $\lambda x.\, x\ x$ is compactified into $\lambda x.\, \mathsf{let}\ f = x\ \mathsf{in}\ f\ f$. Such explicit sharing of abstraction variables is excessive for many applications. This is easily resolved, by unsharing variable vertices before applying the readback, or by preventing the readback from introducing function bindings when only a variable vertex is shared.

§ 3.10.2 (avoiding aliases produced by the readback)**.** The readback function in section 3.7 is sensitive to the degree of sharing of S-vertices in the given λ-term-graph: it maps two λ-term-graphs that only differ in what concerns sharing of S-vertices to different $\lambda_{\mathsf{letrec}}$-terms. Typically, for λ-term-graphs with maximal sharing of S-vertices this can produce function bindings that are just

'aliases', such as g is alias for I in L_3' from example 3.6.18. This can be avoided in two ways: by slightly adapting the readback function, or by performing maximal unsharing of S-vertices before applying the readback as defined.

§ 3.10.3 (preventing disadvantageous sharing)**.** Introducing sharing at compile-time can cause 'space leaks', i.e. a needlessly high memory footprint, at run-time, because 'a large data structure becomes shared [...], and therefore its space which before was reclaimed by garbage collection now cannot be reclaimed until its last reference is used' [50]. For this reason, realisations of CSE [11] prevent the introduction of such undesired sharing by suitable conditions that account for the type of potentially shared subexpressions, and their strictness in the program.

As our approach generalises CSE, it inherits this weakness, and the introduction sharing needs to be restricted in a similar fashion. Technically this can easily achieved by adding additional backlinks to prevent parts of the λ-term-graph from collapsing.

§ 3.10.4 (a more general notion of readback)**.** property (P3) is rather rigorous in that it imposes sharing structure on $\boldsymbol{\lambda}_{\mathsf{letrec}}$ that is specific to λ-term-graphs (degrees of S-sharing). For a weaker definition of property (P3) with $\leftrightarrows^{\mathsf{S}}$ in place of isomorphism, a readback does not have to be injective and independently of the degree of S-sharing a readback function always exists.

§ 3.10.5 (scope-closure strategies)**.** We focused on eager-scope translations since they maximise sharing. However, every scope-closure strategy [19] induces a translation and its own notion of maximal sharing.

Extensions

§ 3.10.6 (full functional languages)**.** In order to support programming languages that are based on $\boldsymbol{\lambda}_{\mathsf{letrec}}$ like Haskell, additional language constructs need to be supported. Such languages can typically be desugared into a core language, which comprises only a small subset of language constructs such as constructors, case statements, and primitives. These constructs can be represented in an extension of $\boldsymbol{\lambda}_{\mathsf{letrec}}$ by additional function symbols. In conjunction with a desugarer our methods are applicable to full programming languages.

§ 3.10.7 (other programming languages, and calculi with binding constructs)**.** Most programming languages feature constructs for grouping definitions that are similar to letrec. We therefore expect that our methods can be adapted to

many imperative languages in particular, and may turn out to be fruitful for optimising compilers. Our methods for achieving maximal sharing certainly generalise to theoretical frameworks and calculi with binding constructs, such as the π-calculus [41], and higher-order rewrite systems (e.g. CRSs and HRSs, [51]) as used here for the formalisation of $\boldsymbol{\lambda}_{\mathsf{letrec}}$.

§ 3.10.8 (fully-lazy λ-lifting)**.** There is a close connection between our methods and fully-lazy λ-lifting [30, 33]. In particular, the required-variable and scope analysis of a $\lambda_{\mathsf{letrec}}$-term L on which the λ-term-graph translation $[\![L]\!]_{\mathcal{T}}$ is based is analogous to the one needed for extracting from L the supercombinators in the result $\hat{L}$ of fully-lazy λ-lifting L. Moreover, the fully-lazy λ-lifting transformation can even be implemented in a natural way on the basis of our methods. Namely as the composition $rb_{LL} \circ [\![\cdot]\!]_{\mathcal{T}}$ of the translation $[\![\cdot]\!]_{\mathcal{T}}$ into λ-term-graph, where rb_{LL} is a variant readback function that, for a given λ-term-graph, synthesises the system $\hat{L}$ of supercombinators, instead of the $\lambda_{\mathsf{letrec}}$-term $rb(L)$.

§ 3.10.9 (maximal sharing on supercombinator translations of $\lambda_{\mathsf{letrec}}$-terms)**.** $\lambda_{\mathsf{letrec}}$-terms L correspond to supercombinator systems S, the result of fully-lazy lambda-lifting L: the combinators in S correspond with 'extended scopes' in L, and supercombinator reduction steps on S correspond with weak β-reduction steps L. In the case of λ-calculus this has been established by Balabonski [5]. Via this correspondence the maximal-sharing method for $\lambda_{\mathsf{letrec}}$-terms can be lifted to obtain a maximal-sharing method systems of supercombinators obtained by fully-lazy lambda-lifting.

§ 3.10.10 (non-eager scope-closure strategies)**.** We focused on eager-scope translations, because they facilitate maximal sharing, and guarantee that interpretations of unfolding-equivalent $\lambda_{\mathsf{letrec}}$-terms are bisimilar. Yet every scope-closure strategy induces a translation and its own notion of maximal sharing. For adapting our maximal sharing method it however necessary to modify the translation into first-order term graphs in such a way that the image class obtained is closed under functional bisimulation ($\mathcal{T}_{1,2}$ is not closed under functional bisimulation, unlike its subclass ${}^{\mathrm{eag}}\mathcal{T}_{1,2} = \mathcal{T}$). This can be achieved by using delimiter vertices also below variable vertices to close scopes that are still open [24, report].

§ 3.10.11 (weaker notions of sharing)**.** The presented methods deal with sharing as expressed by letrec that is horizontal, vertical, or twisted (see remark 3.3.4). By contrast, the μ-construct [8, 20] expresses only vertical, and the non-recursive

let only horizontal, sharing (see remark 3.3.4). By restricting bisimulation (either artificially or by adding special backlinks), our methods can be adapted to the λ-calculus with μ [20], or with let.

Applications

§ 3.10.12 (maximal sharing at run-time)**.** Maximal sharing can be applied repeatedly at run-time in order to regain a maximally shared form, thereby speeding up evaluation. This is reminiscent of 'collapsed tree rewriting' [49] for evaluating first-order term graphs represented as maximally shared dags. Since the state of a program in the memory at run-time is typically represented as a supercombinator graph, compactification by bisimulation collapse can take place directly on that graph (see § 3.10.5), no translation is needed. Compactification can be coupled with garbage collection as bisimulation collapse subsumes some of the work required for a mark and sweep garbage collector. However, a compromise needs to be found between the costs for the optimisation and the gained efficiency.

§ 3.10.13 (additional prevention of disadvantageous sharing)**.** While static analysis methods for preventing sharing that may be disadvantageous at run-time can be adapted from CSE to the maximal-sharing method (see section 3.10), this has yet to be investigated for binding-time analysis [48] and a sharing analysis of partial applications [16]. for fine-tuning sharing of partial applications in supercombinator translations [16].

§ 3.10.14 (compile-time optimisation)**.** Increasing sharing facilitates potential gains in efficiency. Our method generalises common subexpression elimination, but therefore it also inherits its shortcomings: the cost of sharing (e.g. of very small functions) might exceed the gain. In non-strict functional languages, sharing can cause 'memory leaks' [11]. Therefore, similar as for CSE, additional dynamic analyses like binding-time analysis [48], and heuristics to restrict sharing in cases when it is disadvantageous [33, 16] are in order.

§ 3.10.15 (code improvement)**.** In programming it is generally desirable to avoid duplication of code. As an extension of CSE, our method is able to detect code duplication. The bisimulation collapse of the term graph interpretation of a program can, together with the readback, provide guidance on how code can be written more compactly with less duplication. This optimisation has to be fine-tuned to avoid excessive behaviour like the explicit sharing of variable occurrences (see section 3.10). Yet for this only lightweight additional machinery

is needed, such as size constraints or annotations to restrict the bisimulation collapse.

§ 3.10.16 (function equivalence)**.** Recognising whether two programs are equivalent in the sense that they implement the same function is undecidable. Still, this problem is tackled by proof assistants, and by automated theorem provers used in type-checkers of compilers for dependently-typed programming languages such as Agda. For such systems coinductive proofs are more difficult to find than inductive ones, and require more effort by the user. Our method for deciding unfolding-equivalence could contribute to finding coinductive proofs.

Bibliography

[1] Zena M. Ariola and Stefan Blom. Cyclic Lambda Calculi. In *Theoretical Aspects of Computer Software*, volume 1281 of *Lecture Notes in Computer Science*, pages 77–106. Springer Berlin Heidelberg, 1997. doi:10.1007/BFb0014548.

[2] Zena M. Ariola and Jan Willem Klop. Equational Term Graph Rewriting. *Fundamenta Informaticae*, 26(3):207–240, 1996. doi:10.3233/FI-1996-263401, extended version as Vrije Universiteit Amsterdam Technical Report IR–391.

[3] Zena M. Ariola and Jan Willem Klop. Lambda Calculus with Explicit Recursion. *Information and Computation*, 139(2):154–233, 1997. doi:10.1006/inco.1997.2651.

[4] Andrea Asperti and Stefano Guerrini. *The Optimal Implementation of Functional Programming Languages*. Cambridge University Press, 1998.

[5] Thibaut Balabonski. A Unified Approach to Fully Lazy Sharing. In *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 469–480. ACM, 2012. doi:10.1145/2103656.2103713.

[6] Henk P. Barendregt. *The lambda calculus: its syntax and semantics*. Studies in logic and the foundations of mathematics. North-Holland, 1984. http://mathgate.info/cebrown/notes/barendregt.php.

[7] Richard S. Bird and Ross Paterson. de Bruijn notation as a nested datatype. *Journal of Functional Programming*, 9:77–91, 1999. doi:10.1017/S0956796899003366.

[8] Stefan Blom. *Term Graph Rewriting – Syntax and Semantics*. PhD thesis, Vrije Universiteit Amsterdam, 2001.

[9] Mark van den Brand and Paul Klint. ATerms for manipulation and exchange of structured data: It's all about sharing. *Information and Software Technology*, 49(1):55–64, 2007. doi:10.1016/j.infsof.2006.08.009.

[10] N. G. de Bruijn. Lambda Calculus Notation with Nameless Dummies, a Tool for Automatic Formula Manipulation, with Application to the Church-Rosser Theorem. *Indagationes Mathematicae*, 34:381–392, 1972.

[11] Olaf Chitil. Common Subexpressions Are Uncommon in Lazy Functional Languages. In *Selected Papers from the 9th International Workshop (IFL'97)*, Implementation of Functional Languages, pages 53–71. Springer Berlin Heidelberg, 1998. doi:10.1007/BFb0055424.

[12] Alonzo Church. *The Calculi of Lambda-Conversion*. Number 6 in Annals of Mathematical Studies. Princeton University Press, 1941.

[13] Bruno Courcelle. Fundamental properties of infinite trees. *Theoretical Computer Science*, 25(2):95–169, 1983. doi:10.1016/0304-3975(83)90059-2.

[14] Olivier Danvy and Ulrik Pagh Schultz. Lambda-Lifting in Quadratic Time. *Journal of Functional and Logic Programming*, 2004, 2004. doi:10.1007/978-3-540-85373-2_3.

[15] Jörg Endrullis, Clemens Grabmayer, Jan Willem Klop, and Vincent van Oostrom. On Equal μ-Terms. In *Festschrift in Honour of Jan Bergstra*, Special Issue of TCS, 412 (28), pages 3175–3202. Elsevier, 2011.

[16] Benjamin Goldberg and Paul Hudak. Detecting Sharing of Partial Applications in Functional Programs. Technical Report YALEU/DCS/RR-526, Department of Computer Science, Yale University, 1987.

[17] Clemens Grabmayer and Jan Rochel. Confluent Let-Floating. In *Proceedings of IWC 2013 (2nd International Workshop on Confluence)*, 2013.

[18] Clemens Grabmayer and Jan Rochel. Confluent unfolding in the λ-calculus with letrec. In *Proceedings of IWC 2013 (2nd International Workshop on Confluence)*, 2013.

[19] Clemens Grabmayer and Jan Rochel. Expressibility in the Lambda Calculus with letrec. Technical report, Universiteit Utrecht, 2012. `arXiv:1208.2383`, *Note:* generalises results from [20].

[20] Clemens Grabmayer and Jan Rochel. Expressibility in the Lambda Calculus with μ. Technical report, Universiteit Utrecht, 2013. `arXiv:1304.6284`, *Note:* extended version of [21].

[21] Clemens Grabmayer and Jan Rochel. Expressibility in the Lambda Calculus with μ. In *24th International Conference on Rewriting Techniques and Applications, RTA 2013, June 24-26, 2013, Eindhoven, The Netherlands*, volume 21 of *LIPIcs*, pages 206–222. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013. `doi:10.4230/LIPIcs.RTA.2013.206`.

[22] Clemens Grabmayer and Jan Rochel. Maximal Sharing in the Lambda Calculus with letrec. Technical report, Vrije Universiteit Amsterdam and Universiteit Utrecht, 2014. `arXiv:1401.1460`, *Note:* extended version of [23].

[23] Clemens Grabmayer and Jan Rochel. Maximal sharing in the Lambda calculus with letrec. In *Proceedings of the 19th ACM SIGPLAN international conference on Functional programming, Gothenburg, Sweden, September 1-3, 2014*, pages 67–80. ACM, 2014. `doi:10.1145/2628136.2628148`.

[24] Clemens Grabmayer and Jan Rochel. Term Graph Representations for Cyclic Lambda-Terms. Technical report, Universiteit Utrecht, 2013. `arXiv:1308.1034`, *Note:* extended version of [25].

[25] Clemens Grabmayer and Jan Rochel. Term Graph Representations for Cyclic Lambda-Terms. In *Proceedings 7th International Workshop on Computing with Terms and Graphs, TERMGRAPH 2013, Rome, Italy, 23th March 2013.*, number 110 in EPTCS, pages 56–73, 2013. `doi:10.4204/EPTCS.110.7`.

[26] Dimitri Hendriks and Vincent van Oostrom. λ. In *Automated Deduction – CADE-19*, volume 2741 of *Lecture Notes in Computer Science*, pages 136–150. Springer Berlin Heidelberg, 2003. `doi:10.1007/978-3-540-45085-6_11`.

[27] J. E. Hopcroft. An $n \log n$ Algorithm for Minimizing States in a Finite Automata. In *Proceedings of an International Symposium on the Theory*

of Machines and Computations, Haifa, Isreal, pages 189–196. Academic Press, 1971.

[28] J. E. Hopcroft and R. E. Karp. A Linear Algorithm for Testing Equivalence of Finite Automata. Technical report, Cornell University, 1971.

[29] John Hughes. Graph Reduction with Supercombinators. Technical Report PRG28, Oxford University Computing Laboratory, 1982. `http://www.cs.ox.ac.uk/files/3294/PRG28.pdf`.

[30] R. J. M. Hughes. Supercombinators: A new implementation method for applicative languages. In *LFP '82: Proceedings of the 1982 ACM symposium on LISP and functional programming*, pages 1–10, 1982. *Note:* see also [29].

[31] Jeroen Ketema and Jakob Grue Simonsen. Infinitary Combinatory Reduction Systems. *Information and Computation*, 209(6):893 – 926, 2011. `doi:10.1016/j.ic.2011.01.007`.

[32] Thomas Johnsson. Lambda lifting: Transforming programs to recursive equations. In *Functional Programming Languages and Computer Architecture*, volume 201 of *Lecture Notes in Computer Science*, pages 190–203. Springer Berlin Heidelberg, 1985. `doi:10.1007/3-540-15975-4_37`.

[33] Simon L. Peyton Jones. *The Implementation of Functional Programming Languages*. Prentice-Hall International Series in Computer Science. Prentice-Hall, Inc., 1987.

[34] Jeroen Ketema and Jakob Grue Simonsen. Infinitary combinatory reduction systems: confluence. *Logical Methods in Computer Science*, 5(4:3):1–29, 2009. `doi:10.2168/LMCS-5(4:3)2009`.

[35] Jeroen Ketema and Jakob Grue Simonsen. Infinitary Combinatory Reduction Systems: Normalising Reduction Strategies. *Logical Methods in Computer Science*, 6(1:7):1–35, 2010. `doi:10.1145/2528934`.

[36] Jan Willem Klop. *Combinatory Reduction Systems*. PhD thesis, Universiteit Utrecht, 1980.

[37] Jan Willem Klop, Vincent van Oostrom, and Femke van Raamsdonk. Combinatory reduction systems: introduction and survey. *Theoretical Computer Science*, 121(1-2):279 – 308, 1993. `doi:10.1016/0304-3975(93)90091-7`.

[38] Dénes Kőnig. *Theorie der Endlichen und Unendlichen Graphen*. AMS Chelsea Publishing, 2001.

[39] Jean-Jacques Lévy. *Réductions correctes et optimales dans le λ-calcul*. PhD thesis, Université, Paris VII, 1978.

[40] P.-A. Melliès. *Description Abstraite des Systèmes de Réécriture*. PhD thesis, l'Université Paris 7, 1996.

[41] Robin Milner. *Communicating and mobile systems: the π-calculus*. Cambridge University Press, 1999.

[42] Marco T. Morazán and Ulrik Pagh Schultz. Optimal Lambda Lifting in Quadratic Time. In *19th International Workshop on Implementation and Application of Functional Languages*, number 5083 in LNCS. Springer Berlin Heidelberg, 2008. doi:10.1007/978-3-540-85373-2_3.

[43] Daphne A. Norton. Algorithms for Testing Equivalence of Finite Automata, with a Grading Tool for JFLAP. Master's thesis, Department of Computer Science, Rochester Institute of Technology, 2009. http://scholarworks.rit.edu/theses/6939/.

[44] Enno Ohlebusch. Church-Rosser theorems for abstract reduction modulo an equivalence relation. In *Rewriting Techniques and Applications*, volume 1379 of *Lecture Notes in Computer Science*, pages 17–31. Springer Berlin Heidelberg, 1998. doi:10.1007/BFb0052358.

[45] Vincent van Oostrom. *Confluence for Abstract and Higher-Order Rewriting*. PhD thesis, Vrije Universiteit Amsterdam, 1994.

[46] Vincent van Oostrom. FD à la Melliès, 1997. http://www.phil.uu.nl/~oostrom/publication/ps/FDalaMellies.ps, Vrije Universiteit Amsterdam.

[47] Vincent van Oostrom, Kees-Jan van de Looij, and Marijn Zwitserlood. Lambdascope. Extended Abstract for the Workshop on Algebra and Logic on Programming Systems (ALPS), Kyoto, April 10th 2004, 2004. http://www.phil.uu.nl/~oostrom/publication/pdf/lambdascope.pdf.

[48] J. Palsberg and M. I. Schwartzbach. Binding-time analysis: abstract interpretation versus type inference. In *International Conference on Computer Languages*, pages 289–298, 1994. doi:10.1109/ICCL.1994.288372.

[49] Detlef Plump. *Evaluation of Functional Expressions by Hypergraph Rewriting*. PhD thesis, Universität Bremen, 1993.

[50] André Luís de Medeiros Santos. *Compilation by Transformation in Non-Strict Functional Languages*. PhD thesis, University of Glasgow, 1995.

[51] TeReSe. *Term Rewriting Systems*, volume 55 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2003. `https://www.cs.vu.nl/~terese/`.

[52] Michael Jonathan Thyer. *Lazy specialization*. PhD thesis, University of York, 2003.

[53] Christopher Peter Wadsworth. *Semantics and pragmatics of the lambda-calculus*. PhD thesis, University of Oxford, 1971.

Appendix A

Examples: Graph Translation

§ A.17 (overview)**.** For two terms from chapter 3 we provide the fully worked out stepwise translation of $\lambda_{\mathsf{letrec}}$-terms according to $\mathcal{R}$ from fig. 3.3 producing λ-ap-ho-term-graphs and for the rules $\mathcal{R}_{\mathsf{S}}$ from fig. 3.7 producing λ-term-graphs. Note that in each step of the translation we apply not only one but multiple translation rules (indicated as subscripts of the steps), one for each translation box of the current graph. When no more rules are applicable, indirection vertices are erased.

§ A.18. We start off by a simple example, namely the translation (fig. 3.2) of the term P from example 3.3.3. There is no application of the S-rule, thus it yields the exact same translation sequence for $\mathcal{R}$ as for $\mathcal{R}_{\mathsf{S}}$ producing a λ-term-graph.

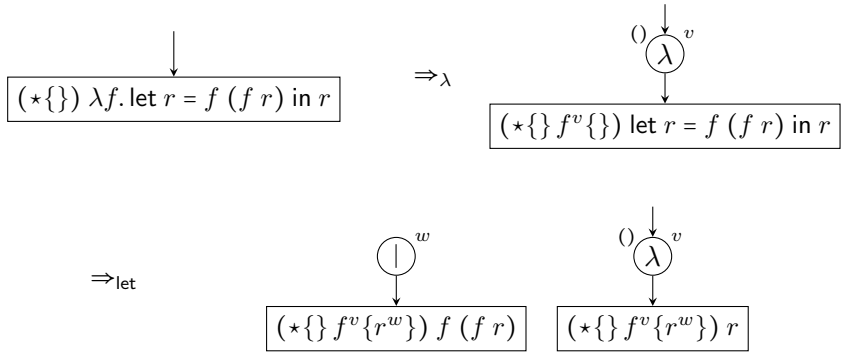

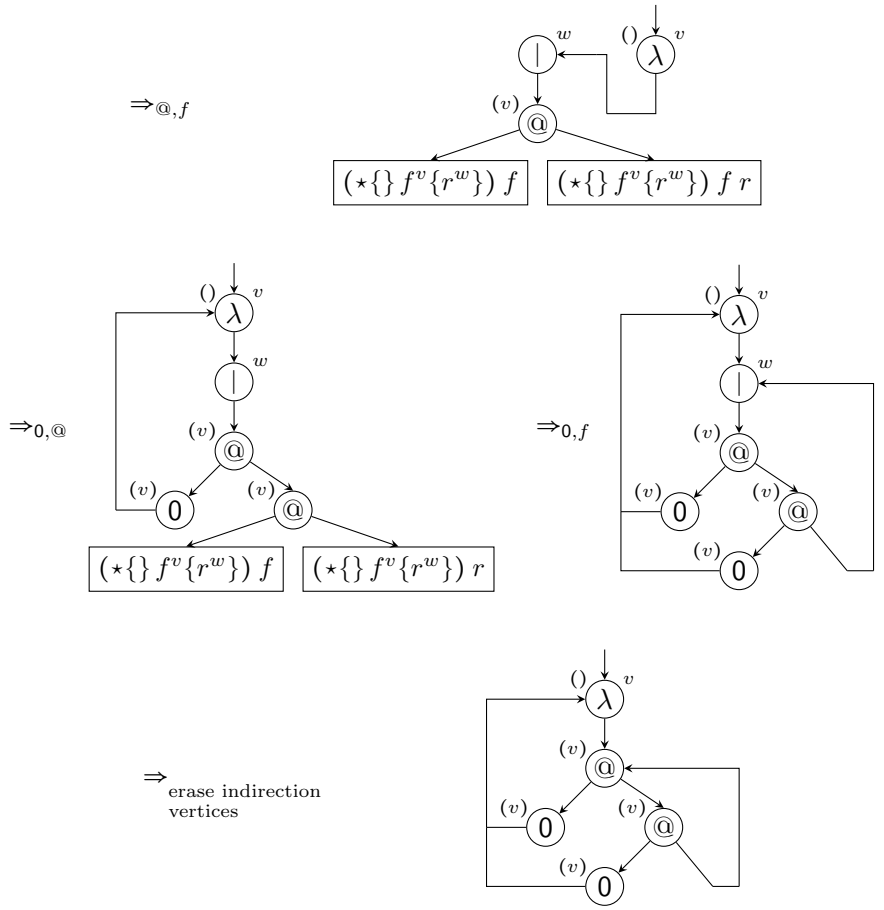

§ A.19. We continue with the term L_2 from example 3.6.18. We translate it in two different ways correspoding to the first-order term graph semantics $[\![\cdot]\!]^{\min}_{\mathcal{T}}$ from definition 3.6.9 and $[\![\cdot]\!]_{\mathcal{T}}$ from definition 3.6.15, respectively. Both sequences of translation steps yield the same λ-ap-ho-term-graph $[\![L_2]\!]_{\mathcal{H}}$,[1] but we obtain two different λ-term-graphs $[\![L_2]\!]^{\min}_{\mathcal{T}}$ and $[\![L_2]\!]_{\mathcal{T}}$ as in fig. 3.9.

§ A.20 (dotted S-vertices)**.** Since we want to illustrate translations w.r.t. $\mathcal{R}$ which produces λ-ap-ho-term-graphs (which do not contain S-vertices) and w.r.t.

[1] This is due to theorem 2.7.20, proposition 3.6.11, and proposition 3.6.16.

$\mathcal{R}_{\mathsf{S}}$ which produces λ-term-graphs (which contain S-vertices) at the same time, in the graphs S-vertices drawn with dotted lines.

§ A.21 (prefix lengths in the let-rule)**.** Both translations are eager-scope (i.e. applications of S-rules are given priority, and prefixes lengths are chosen small enough in the let-rule, see § 3.5.17) but differ in how they resolve the non-determinism due to different possible choices for the prefix lengths $l_1, \ldots, l_k$ in the let-rule.

§ A.22 (translation 1: no S-sharing)**.** First we consider the translations of L_2 with $[\![\cdot]\!]_{\mathcal{H}}$ and $[\![\cdot]\!]_{\mathcal{T}}^{\min}$, i.e. the translation process in which prefix lengths are chosen minimally when applying the let-rule resulting in no sharing of S-vertices.

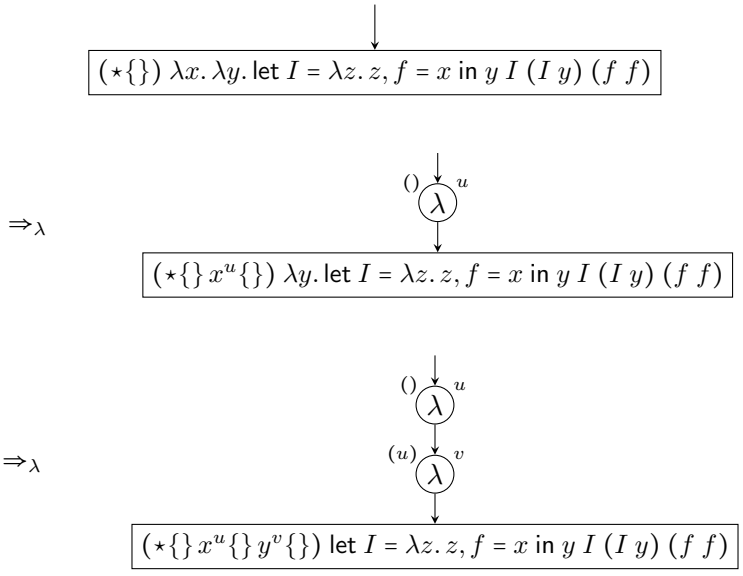

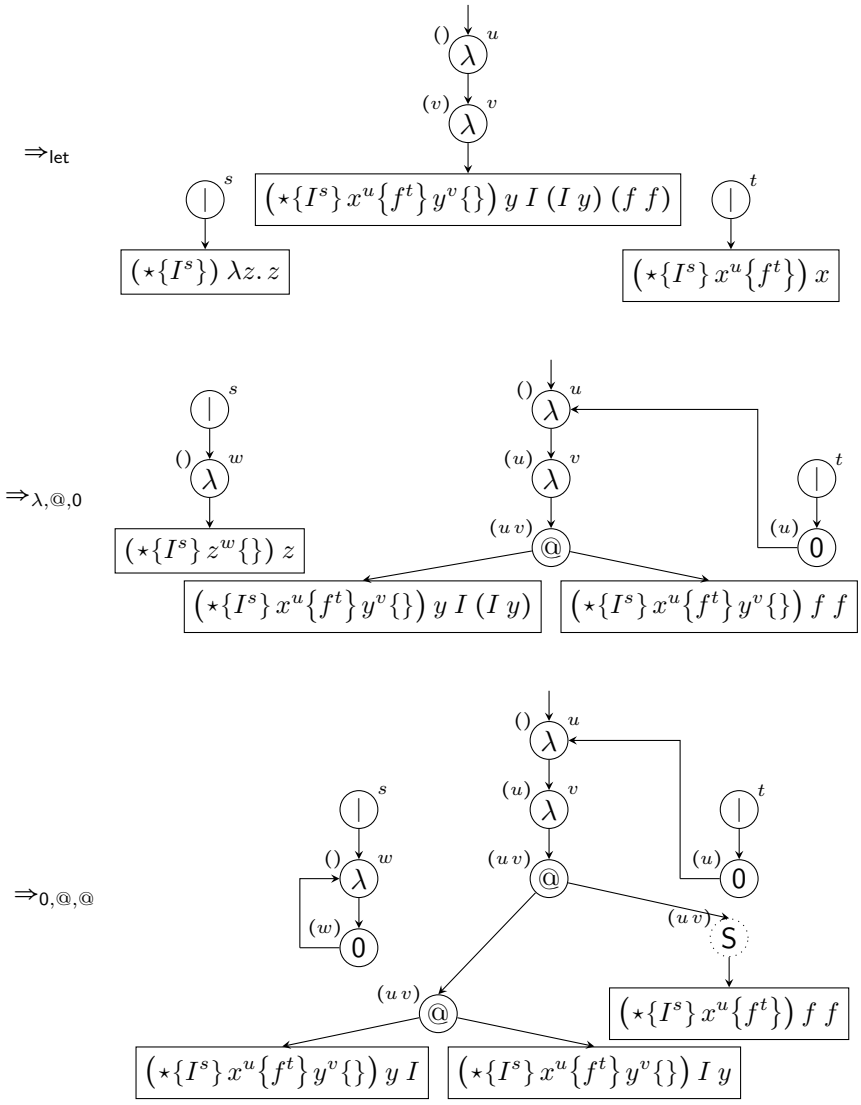
⇒let
⇒λ,@,0
⇒0,@,@
λ
@
0
S

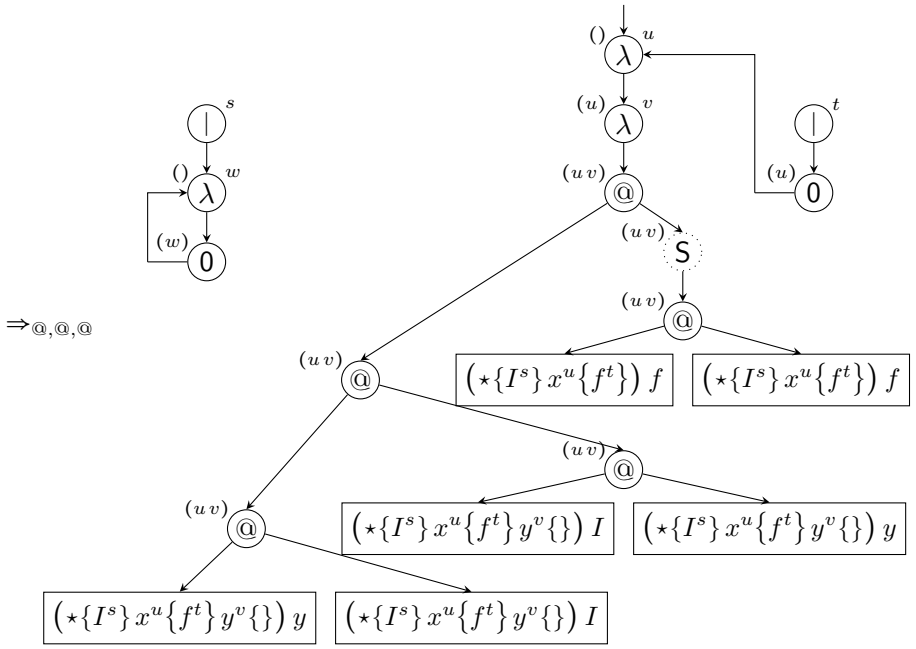

$\Rightarrow_{@,@,@}$
$\left(\star\{I^s\}\, x^u\{f^t\}\right) f$
$\left(\star\{I^s\}\, x^u\{f^t\}\right) f$
$\left(\star\{I^s\}\, x^u\{f^t\}\, y^v\{\}\right) I$
$\left(\star\{I^s\}\, x^u\{f^t\}\, y^v\{\}\right) y$
$\left(\star\{I^s\}\, x^u\{f^t\}\, y^v\{\}\right) y$
$\left(\star\{I^s\}\, x^u\{f^t\}\, y^v\{\}\right) I$

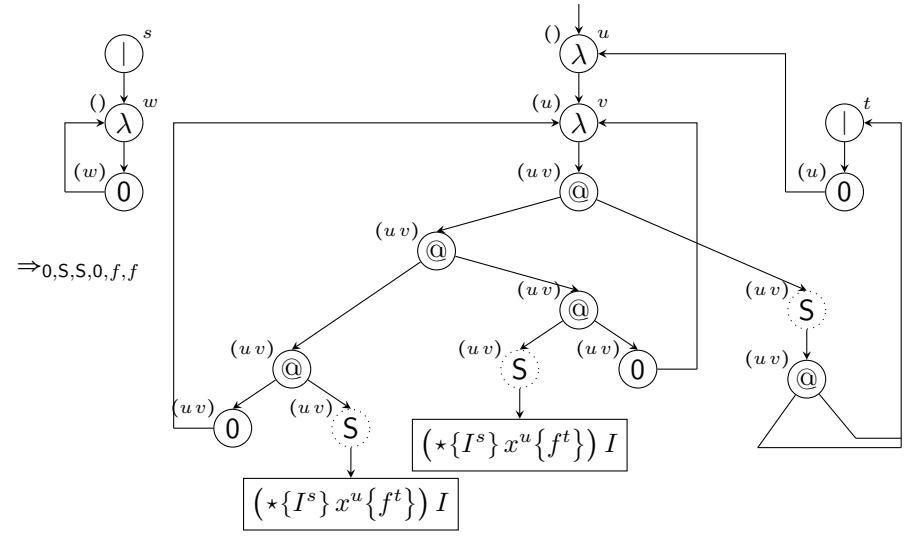

$\Rightarrow_{0,S,S,0,f,f}$
$\left(\star\{I^s\}\, x^u\{f^t\}\right) I$
$\left(\star\{I^s\}\, x^u\{f^t\}\right) I$

$\Rightarrow_{\mathsf{S,S}}$

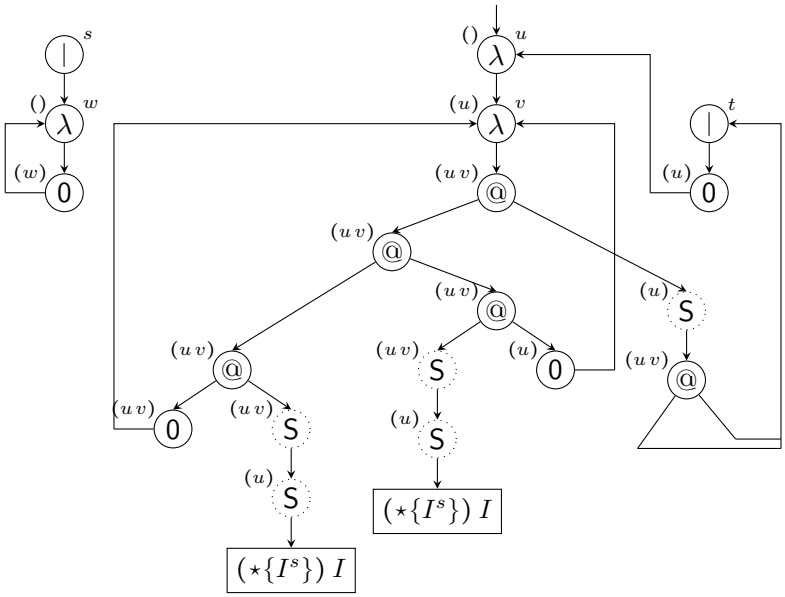

$\Rightarrow_{f,f}$

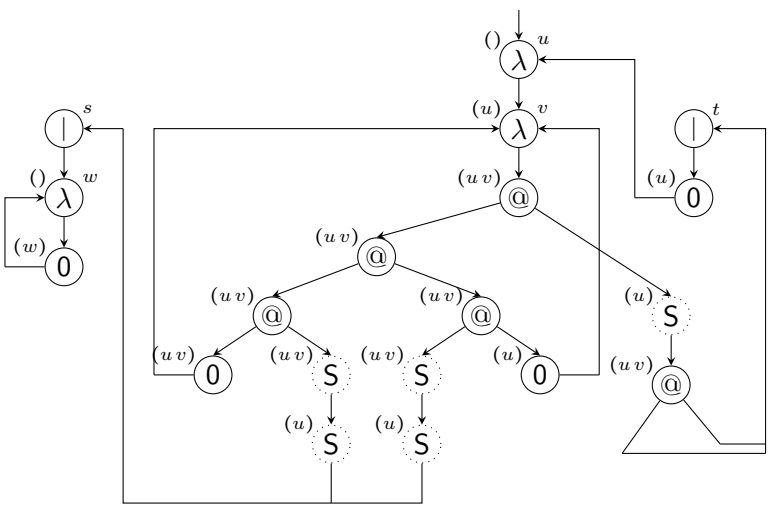

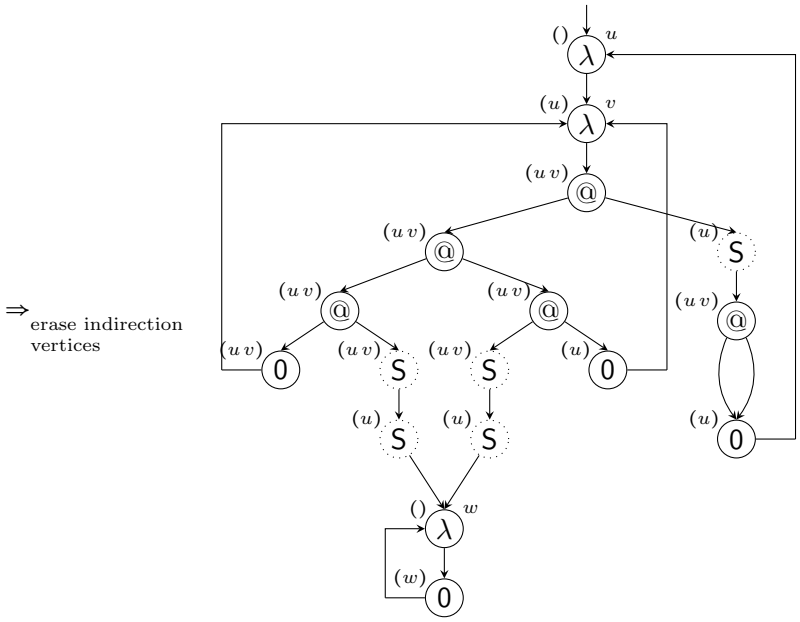

§ A.23 (translation 2: maximal S-sharing)**.** Second, we give the translation of the same term L_2 (from example 3.6.18) with the process needed for the first-order term graph semantics $[\![\cdot]\!]_{\mathcal{T}}$, yielding $[\![L_2]\!]_{\mathcal{T}}$ in fig. 3.3, and the corresponding λ-ap-ho-term-graph $[\![L_2]\!]_{\mathcal{H}}$. The obtained λ-term-graph $[\![L_2]\!]_{\mathcal{T}}$ differs from the λ-term-graph $[\![L_2]\!]_{\mathcal{T}}^{\min}$ obtained in § A.22 by a higher degree of S-sharing.

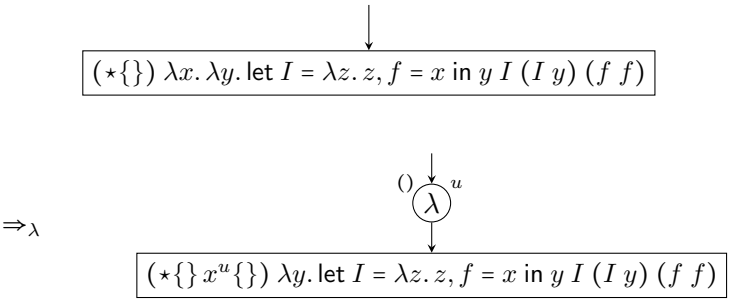

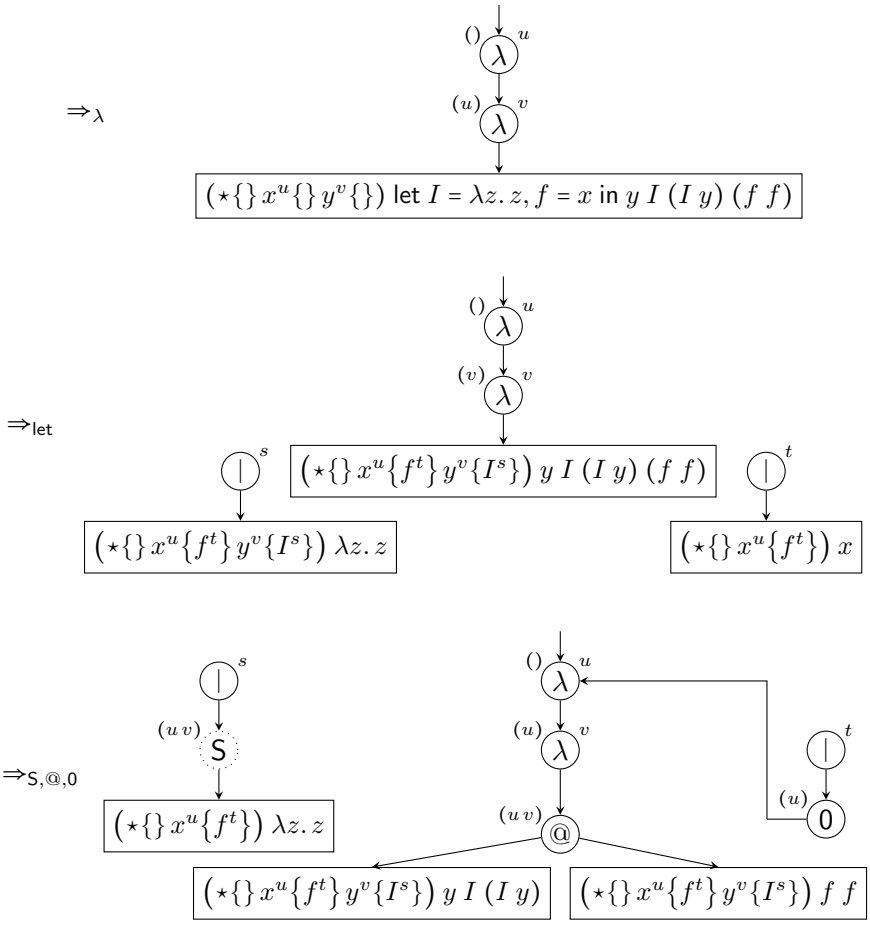
⇒λ
()
u
λ
(u)
v
λ
(⋆{} x^u{} y^v{}) let I = λz. z, f = x in y I (I y) (f f)
⇒let
()
u
λ
(v)
v
λ
s
|
(⋆{} x^u{f^t} y^v{I^s}) y I (I y) (f f)
t
|
(⋆{} x^u{f^t} y^v{I^s}) λz. z
(⋆{} x^u{f^t}) x
⇒S,@,0
s
|
(u v)
S
(⋆{} x^u{f^t}) λz. z
()
u
λ
(u)
v
λ
(u v)
@
t
|
(u)
0
(⋆{} x^u{f^t} y^v{I^s}) y I (I y)
(⋆{} x^u{f^t} y^v{I^s}) f f

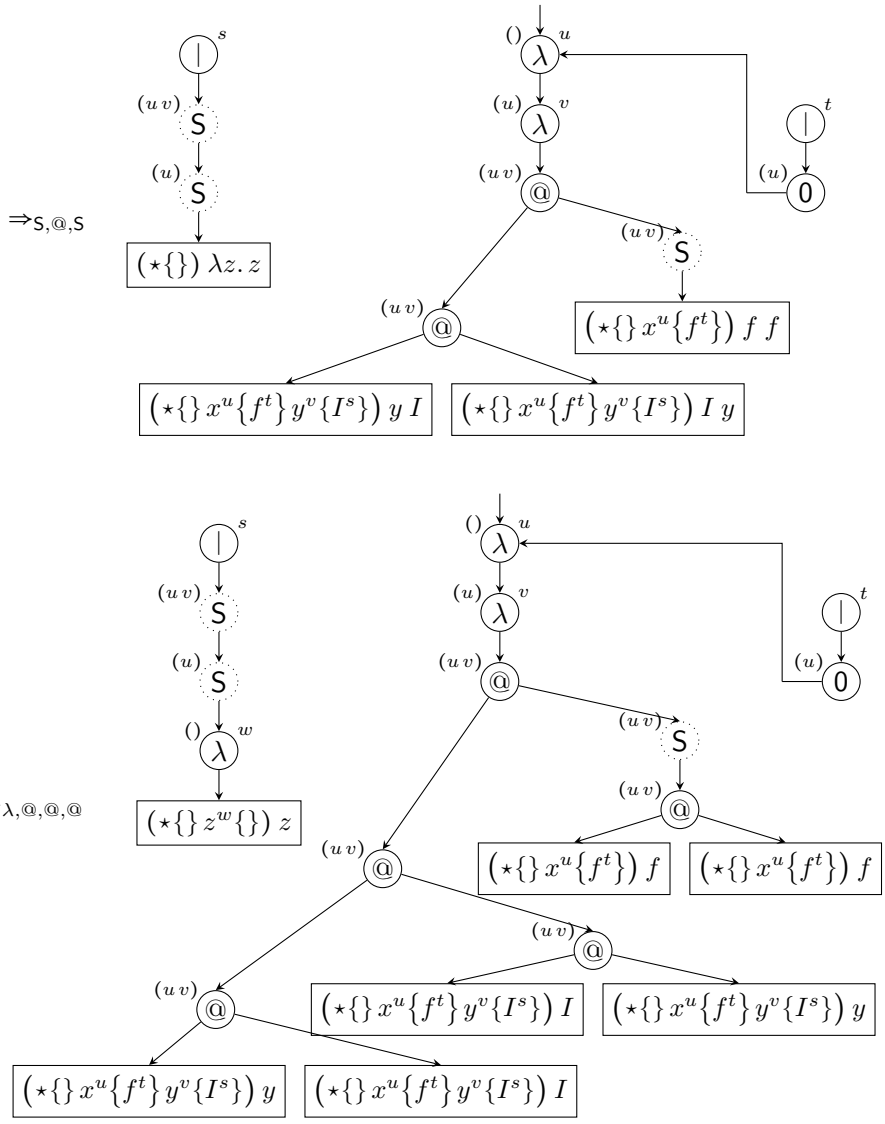
⇒S,@,S
s
(u v)
S
(u)
S
(⋆{}) λz. z
()
u
λ
(u)
v
λ
(u v)
@
t
(u)
0
(u v)
S
(⋆{} x^u{f^t}) f f
(u v)
@
(⋆{} x^u{f^t} y^v{I^s}) y I
(⋆{} x^u{f^t} y^v{I^s}) I y
⇒λ,@,@,@
s
(u v)
S
(u)
S
()
w
λ
(⋆{} z^w{}) z
()
u
λ
(u)
v
λ
(u v)
@
t
(u)
0
(u v)
S
(u v)
@
(⋆{} x^u{f^t}) f
(⋆{} x^u{f^t}) f
(u v)
@
(u v)
@
(⋆{} x^u{f^t} y^v{I^s}) I
(⋆{} x^u{f^t} y^v{I^s}) y
(⋆{} x^u{f^t} y^v{I^s}) y
(⋆{} x^u{f^t} y^v{I^s}) I

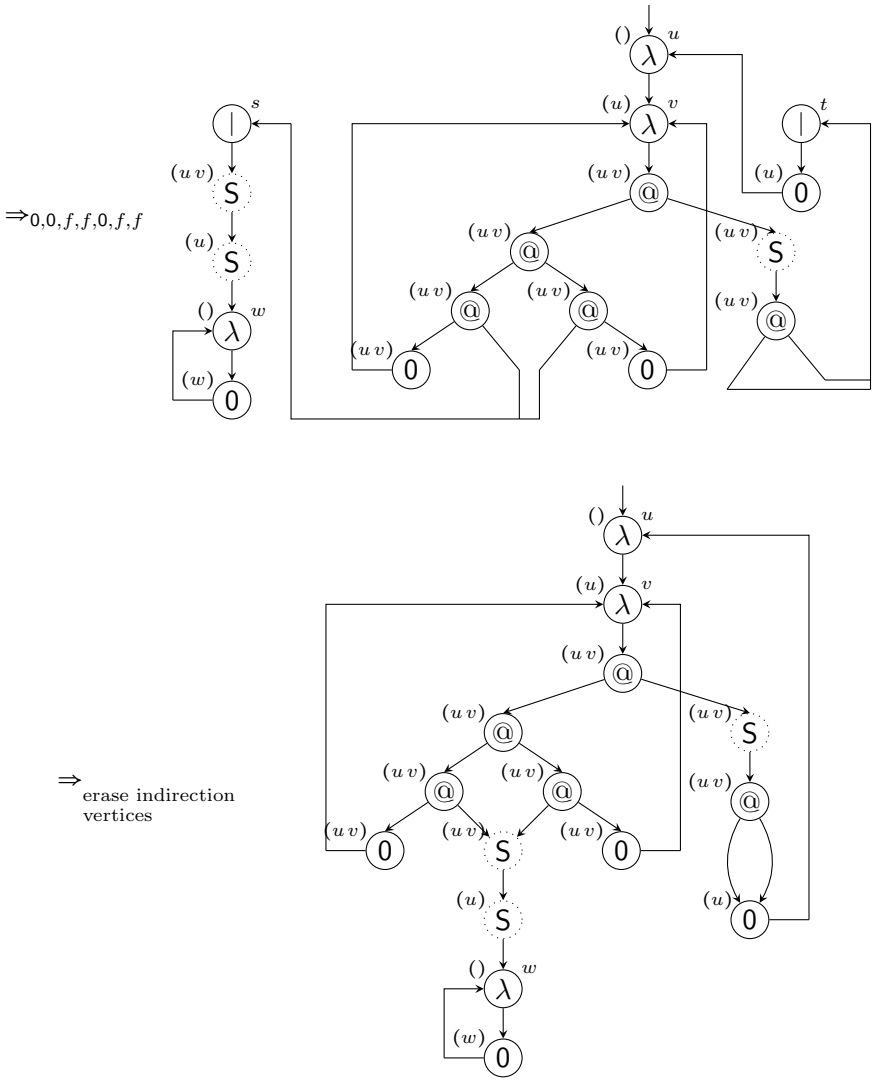
⇒0,0,f,f,0,f,f
s
t
u
v
w
()
(u)
(w)
(u v)
λ
@
S
0
⇒
erase indirection
vertices

Appendix B

Implementation Showcase

To demonstrate the realisability of our method, we have implemented the methods from chapter 3. The implementation is called maxsharing and is available on Hackage: `http://hackage.haskell.org/package/maxsharing/`. It is written in Haskell and therefore requires the Haskell Platform to be installed. Then, maxsharing can be installed via cabal-install using the commands `cabal update` and `cabal install maxsharing` from the terminal. Invoke the executable `maxsharing` in your cabal-directory with a file as an argument that contains a $\lambda_{\mathsf{letrec}}$-term. Run `maxsharing -h` for help on run-time flags.

For further illustration we provide the output of the maxsharing tool for a number of examples from the thesis. Since the volume of the output is too large to include here it is supplied as an external appendix available at `http://rochel.info/thesis/`.

Appendix C

Confluence of unfolding $\lambda_{\mathsf{letrec}}$-terms

Here we present a proof of proposition 0.7.8 which is an adaptation of [18] where we prove the confluence of a precursor of $\boldsymbol{R}_{\nabla}$. Here the proof below is written out for $\boldsymbol{R}_{\nabla_\bullet}$, but works for $\boldsymbol{R}_{\nabla}$ as well if the critical pairs involving the rules ϱ^{tighten} and $\varrho^{\bullet}$ are left out.

Proof of proposition 0.7.8. First of all, we cannot use Newman's Lemma to prove confluence, because $\boldsymbol{R}_{\nabla_\bullet}$ is not terminating. To show confluence of $\boldsymbol{R}_{\nabla_\bullet}$ we use the method of 'decreasing diagrams' [45, Section 2.3] [51, Section 14.2]. We use it however not to prove confluence of the rewriting relation $\to_{\nabla}$ induced by $\boldsymbol{R}_{\nabla_\bullet}$ directly, but of the abstract reduction system $\mathcal{A} = (\mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}}), \{\twoheadrightarrow_{\rho_d} \mid (d,\rho) \in \mathbb{N} \times R\})$ with R as the set of rules of $\boldsymbol{R}_{\nabla_\bullet}$ where $\twoheadrightarrow_{\rho_d}$ denotes the parallel rewriting relation on $\mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})$ induced by rule ρ at letrec-depth d. As a precedence order we consider the order induced by the letrec-depth:

$$\twoheadrightarrow_{\rho_d} \geq \twoheadrightarrow_{\sigma_{d'}} \iff d \geq d'$$

The letrec-depth of a redex in $\lambda_{\mathsf{letrec}}$-term denotes the number of let-nodes passed on the path from the root of the term tree to the corresponding position. We write $\to_{\rho_d}$ to denote the relation induced by applying rule ρ contracting a redex at letrec-depth d.

Let us denote the rewriting relation induced by $\mathcal{A}$ by $\to_{\mathcal{A}}$:

$$\to_{\mathcal{A}} = \bigcup \{\twoheadrightarrow_{\rho_d} \mid (d,\rho) \in \mathbb{N} \times R\}$$

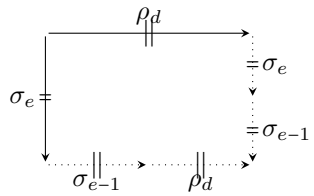

Figure C.1. Elementary diagram

If $\to_{\mathcal{A}}$ is confluent then $\to_{\triangledown}$ is confluent because it holds: $\to_{\triangledown} \subseteq \to_{\mathcal{A}} \subseteq \twoheadrightarrow_{\triangledown}$ or equivalently $\twoheadrightarrow_{\mathcal{A}} = \twoheadrightarrow_{\triangledown}$ (see also [45, Lemma 2.2.5]).

We use parallel steps because the preceding attempt to prove confluence of $\boldsymbol{R}_{\triangledown\bullet}$-steps themselves by decreasing diagrams was unsuccessful. As a precedent order we considered an ordering on the rules and lexicographic extensions of such orderings with the letrec-depth of the contracted redex. We came to the conclusion that no such order could ensure decreasingness of the elementary diagrams of both the critical pairs as well as the strictly nested redexes. This was due to redex duplication induced by the diverging steps, so that joining the diagram required a many-step that disrupted decreasingness. In order to resolve this problem we considered parallel steps such that the problematic multi-step would become a single parallel step. This led to more intricate diagrams but turned out to be a viable solution.

We will show that two diverging parallel steps in $\boldsymbol{R}_{\triangledown\bullet}$ can be joined in an elementary diagram of the following form with $d \leq e$.

If we pick as the precedence order on the steps the order that is induced by their letrec-depth, the diagram is decreasing. Note that in all the diagrams we implicitly assume the reflexive closure for all arrows. The rest of the proof is structured as follows. To justify the diagram we distinguish the cases $d = e$ and $d < e$, for which we construct diagrams that are instances of the diagram in fig. C.1.

Case 1

For $d = e$ we need to consider parallel diverging steps contracting redexes at the same letrec-depth d. We construct the diagram below which is an instance of the diagram above where the diverging parallel steps are in sequentialised form. We write terms as fillings of a multihole context C with all its holes at letrec-depth d such that the contracted ρ_d- and σ_d-redexes are filled into

these holes. In this way we can make explicit at which position a step takes place, i.e. at the root of the context hole fillings. The topmost row and the leftmost column are respective sequentialisations of the parallel diverging ρ_d- and σ_d-steps into single steps.

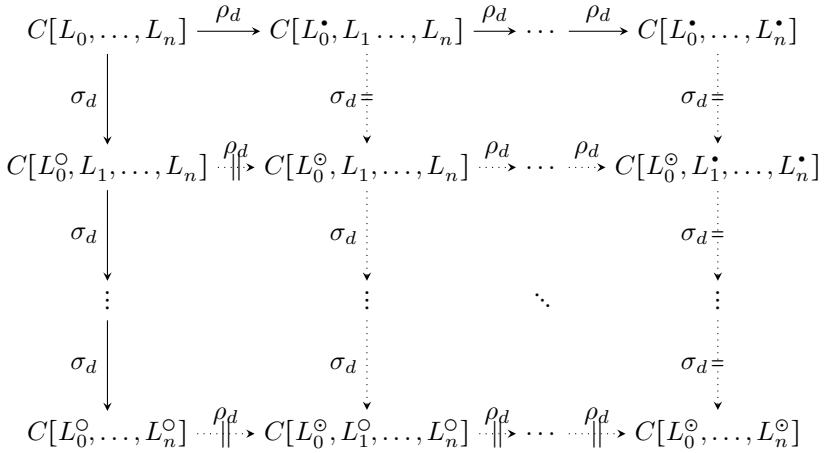

In this diagram the tiles at (i, j) for $i, j \in \{0, \ldots, n-1\}$ below the diagonal $(i < j)$ look as follows:

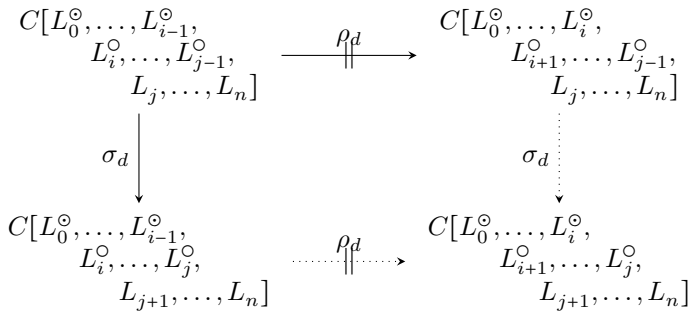

The tiles at (i, j) for $i, j \in \{0, \ldots, n-1\}$ above the diagonal $(i > j)$ look as follows:

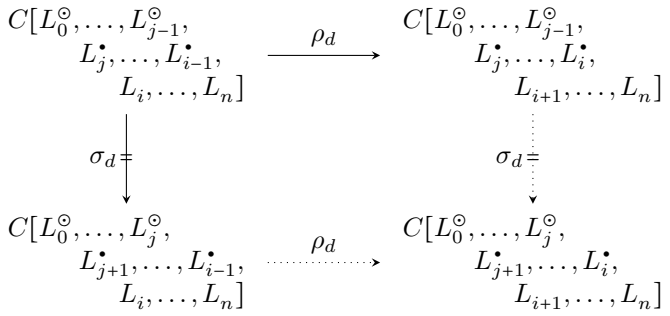

For $i \in \{0,\ldots,n-1\}$ the i-th diagonal tile looks like this:

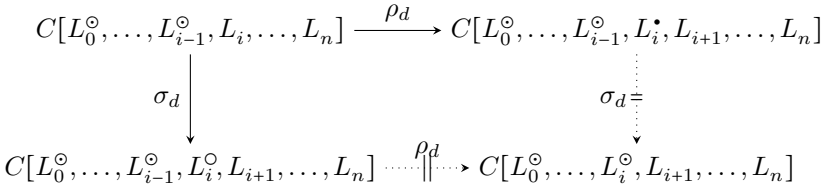

Only the tiles on the diagonal require closer attention because for all other tiles the vertical and horizontal steps take place in different holes of the context, therefore they are disjoint and consequently commute. In the tiles on the diagonal the diverging steps may be either due to a critical pair or to identical steps. In the latter case the diagram is easily joined. In case of a critical pair, since all steps take place at the same letrec-depth any such critical pair must arise from a root overlap. Exhaustive scrutiny of all these critical pairs reveals that they can be joined in a way that conforms to the tiles on the diagonal. Note that the letrec-depths of the steps have to be increased by d according to the lifting into a context with its hole at letrec-depth d.

Remark C.24. Note, that in the critical pair analysis below, we are not always faithful to the actual formulation of the rules, in the sense that we sometimes assume that a binding appears at a specific position in the list of bindings of a let-expression. This is merely to save page space and can be easily generalised.

Critical pairs due to λ_0/ρ_0

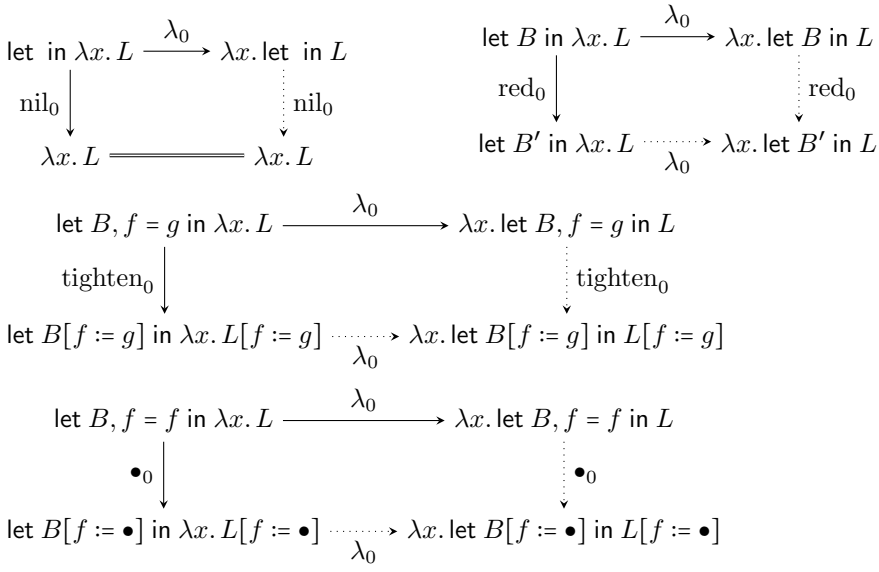

Critical pairs due to $@_0/\sigma_0$

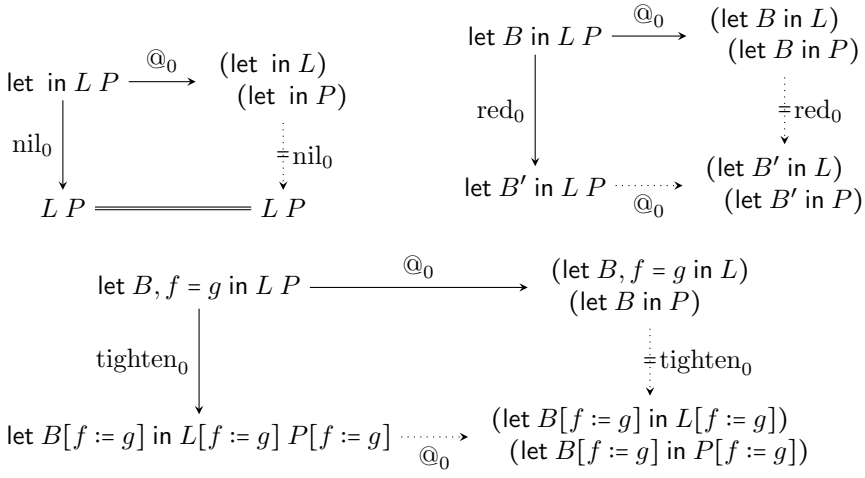

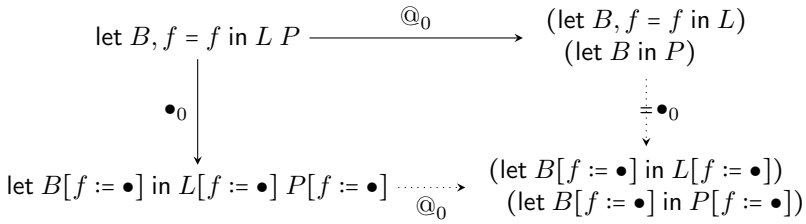

Critical pairs due to $\mathrm{letrec}_0/\sigma_0$

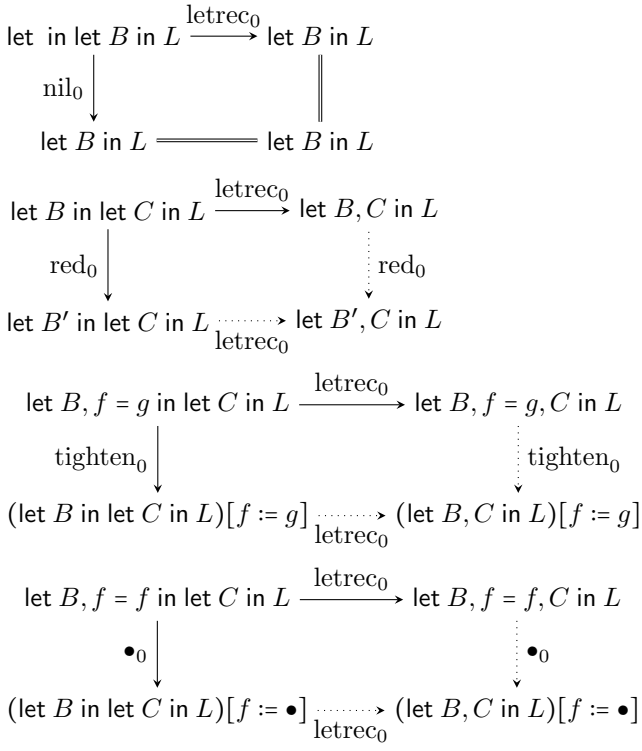

Critical pairs due to rec_0/σ_0

$$\begin{array}{ccc}
\mathsf{let}\ B\ \mathsf{in}\ f_i & \xrightarrow{\mathrm{rec}_0} & \mathsf{let}\ B\ \mathsf{in}\ L_i \\
{\scriptstyle \mathrm{red}_0}\downarrow & & \vdots\,{\scriptstyle \mathrm{red}_0} \\
\mathsf{let}\ B'\ \mathsf{in}\ f_i & \xrightarrow[\mathrm{rec}_0]{} & \mathsf{let}\ B'\ \mathsf{in}\ L_i
\end{array}$$

There is a $\mathrm{rec}_0/\mathrm{tighten}_0$ critical pair. We distinguish the cases, whether the body consists of the 'tightened' variable:

$$\begin{array}{ccc}
\mathsf{let}\ B, f = g\ \mathsf{in}\ f & \xrightarrow{\mathrm{rec}_0} & \mathsf{let}\ B, f = g\ \mathsf{in}\ g \\
{\scriptstyle \mathrm{tighten}_0}\downarrow & & \vdots\,{\scriptstyle \mathrm{tighten}_0} \\
\mathsf{let}\ B[f := g]\ \mathsf{in}\ g & = & \mathsf{let}\ B[f := g]\ \mathsf{in}\ g
\end{array}$$

$$\begin{array}{ccc}
\mathsf{let}\ B, f = g, h = L\ \mathsf{in}\ h & \xrightarrow{\mathrm{rec}_0} & \mathsf{let}\ B, f = g, h = L\ \mathsf{in}\ L \\
{\scriptstyle \mathrm{tighten}_0}\downarrow & & \vdots\,{\scriptstyle \mathrm{tighten}_0} \\
\mathsf{let}\ B[f := g], h = L[f := g]\ \mathsf{in}\ h & \xrightarrow[\mathrm{rec}_0]{} & \mathsf{let}\ B, h = L\ \mathsf{in}\ L[f := g]
\end{array}$$

Likewise, there is a $\mathrm{rec}_0/\bullet_0$ critical pair. We distinguish the cases, whether the body consists of the 'blackholed' variable:

$$\begin{array}{ccc}
\mathsf{let}\ B, f = f\ \mathsf{in}\ f & \xrightarrow{\mathrm{rec}_0} & \mathsf{let}\ B, f = f\ \mathsf{in}\ f \\
{\scriptstyle \bullet_0}\downarrow & & \vdots\,{\scriptstyle \bullet_0} \\
\mathsf{let}\ B[f := \bullet]\ \mathsf{in}\ \bullet & = & \mathsf{let}\ B[f := \bullet]\ \mathsf{in}\ \bullet
\end{array}$$

$$\begin{array}{ccc}
\mathsf{let}\ B, f = f, h = L\ \mathsf{in}\ h & \xrightarrow{\mathrm{rec}_0} & \mathsf{let}\ B, f = f, h = L\ \mathsf{in}\ L \\
{\scriptstyle \bullet_0}\downarrow & & \vdots\,{\scriptstyle \bullet_0} \\
\mathsf{let}\ B[f := \bullet], h = L[f := \bullet]\ \mathsf{in}\ h & \xrightarrow[\mathrm{rec}_0]{} & \mathsf{let}\ B, h = L\ \mathsf{in}\ L[f := \bullet]
\end{array}$$

Critical pairs due to nil_0/σ_0

There are no additional nil_0/σ_0 critical pairs on top of the ones already scrutinised above.

Critical pairs due to red_0/σ_0

There is a $\mathrm{red}_0/\mathrm{tighten}_0$ critical pair, for which we distinguish the cases whether the 'tightened' binding is used.

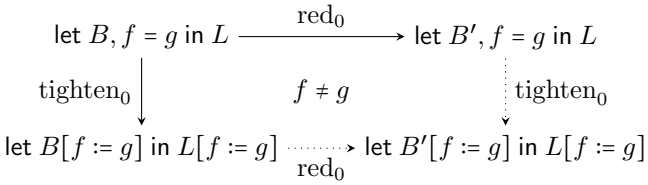

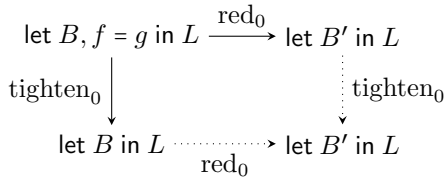

Likewise, there is a $\mathrm{red}_0/\bullet_0$ critical pair, for which we distinguish the cases whether the 'blackholed' binding is used.

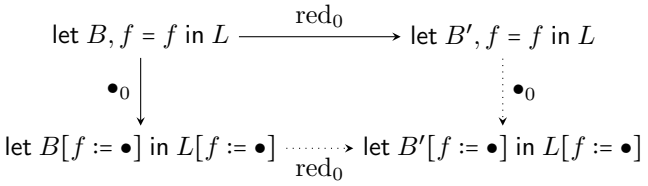

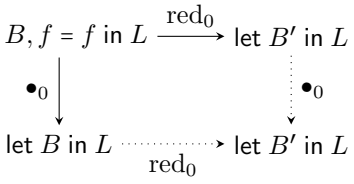

Critical pairs due to $\mathrm{tighten}_0/\sigma_0$

$$\begin{array}{ccc}
\mathsf{let}\ B, f = g, h = i\ \mathsf{in}\ L & \xrightarrow{\mathrm{tighten}_0} & (\mathsf{let}\ B, h = i\ \mathsf{in}\ L)[f := g] \\
\Big\downarrow{\scriptstyle \mathrm{tighten}_0} & f \neq g \quad h \neq i & \vdots\,{\scriptstyle \mathrm{tighten}_0} \\
(\mathsf{let}\ B, f = g\ \mathsf{in}\ L)[h := i] & \xrightarrow[\mathrm{tighten}_0]{} & (\mathsf{let}\ B\ \mathsf{in}\ L)[f := g][h := i]
\end{array}$$

$$\begin{array}{ccc}
\mathsf{let}\ B, f = g, h = h\ \mathsf{in}\ L & \xrightarrow{\mathrm{tighten}_0} & (\mathsf{let}\ B, h = h\ \mathsf{in}\ L)[f := g] \\
\Big\downarrow{\scriptstyle \bullet_0} & f \neq g & \vdots\,{\scriptstyle \bullet_0} \\
(\mathsf{let}\ B, f = g\ \mathsf{in}\ L)[h := \bullet] & \xrightarrow[\mathrm{tighten}_0]{} & (\mathsf{let}\ B\ \mathsf{in}\ L)[f := g][h := \bullet]
\end{array}$$

Critical pairs due to $\bullet_0/\sigma_0$

$$\begin{array}{ccc}
\mathsf{let}\ B, f = f, g = g\ \mathsf{in}\ L & \xrightarrow{\bullet_0} & (\mathsf{let}\ B, g = g\ \mathsf{in}\ L)[f := \bullet] \\
\Big\downarrow{\scriptstyle \bullet_0} & & \vdots\,{\scriptstyle \bullet_0} \\
(\mathsf{let}\ B, f = f\ \mathsf{in}\ L)[g := \bullet] & \xrightarrow[\bullet_0]{} & (\mathsf{let}\ B\ \mathsf{in}\ L)[f := \bullet][g := \bullet]
\end{array}$$

Case 2

For $d < e$ we use the same approach as for $d = e$, the diagram is however more involved. Again, we use a context C with context holes at letrec-depth d. But since $e > d$, more than one σ_e-contraction may take place in one such hole. Therefore a per-hole partitioning of the vertical steps requires a sequence of parallel steps.

The diagram below fits the scheme of the elementary diagram (fig. C.1) when interleaving the σ_e-steps with the σ_{e-1}-steps in the rightmost column such that steps at depth e preceed those at depth $e-1$. Similarly for the bottommost row where the ρ_{e-1}-steps have to preceed the σ_d-steps. These reorderings are possible since the segments represent contractions within different holes of C. As in the previous diagram the tiles which do not lie on the diagonal are

unproblematic, which leaves us to complete the proof by constructing the tiles on the diagonal.

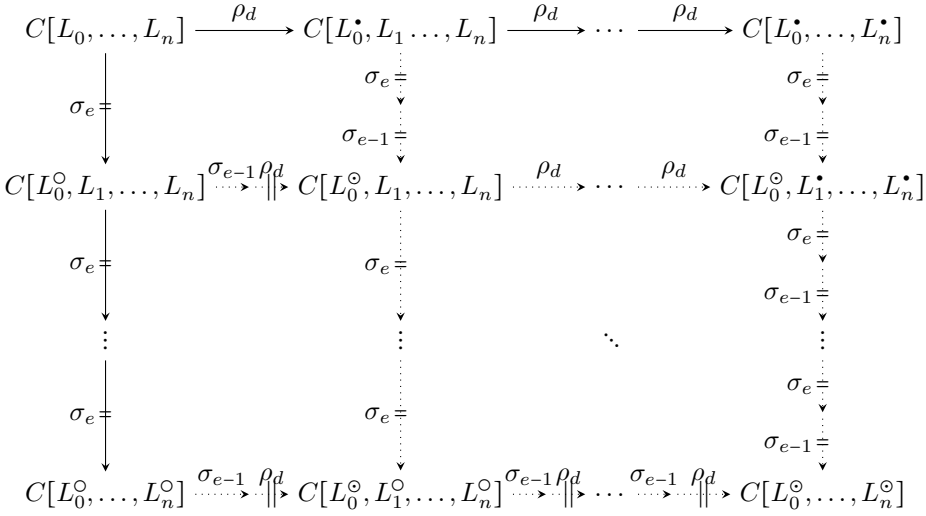

In this diagram the tiles at (i,j) for $i,j \in \{0,\ldots,n-1\}$ below the diagonal $(i<j)$ look as follows:

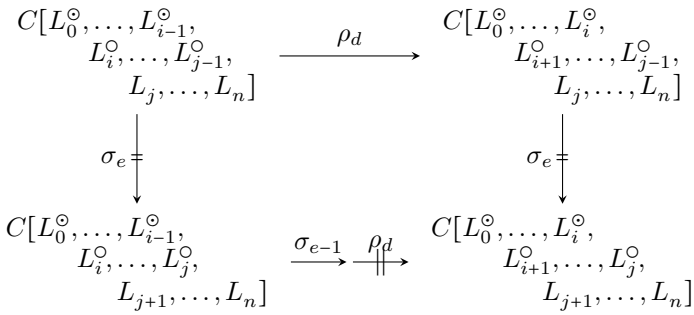

The tiles at (i,j) for $i,j \in \{0,\ldots,n-1\}$ above the diagonal $(i>j)$ look as follows:

$$
\begin{array}{ccc}
\begin{array}{l} C[L_0^{\odot},\dots,L_{i-1}^{\odot}, \\ \quad L_i^{\bullet},\dots,L_{j-1}^{\bullet}, \\ \qquad L_j,\dots,L_n] \end{array} & \xrightarrow{\rho_d} & \begin{array}{l} C[L_0^{\odot},\dots,L_{i-1}^{\odot}, \\ \quad L_i^{\bullet},\dots,L_j^{\bullet}, \\ \qquad L_{j+1},\dots,L_n] \end{array} \\
\sigma_e \Downarrow & & \sigma_e \Downarrow \\
\sigma_{e-1} \Downarrow & & \sigma_{e-1} \Downarrow \\
\begin{array}{l} C[L_0^{\odot},\dots,L_i^{\odot}, \\ \quad L_{i+1}^{\bullet},\dots,L_{j-1}^{\bullet}, \\ \qquad L_j,\dots,L_n] \end{array} & \xrightarrow{\rho_d} & \begin{array}{l} C[L_0^{\odot},\dots,L_i^{\odot}, \\ \quad L_{i+1}^{\bullet},\dots,L_j^{\bullet}, \\ \qquad L_{j+1},\dots,L_n] \end{array}
\end{array}
$$

For $i \in \{0,\dots,n-1\}$ the i-th diagonal tile looks like this:

$$
\begin{array}{ccc}
C[L_0^{\odot},\dots,L_{i-1}^{\odot},L_i,\dots,L_n] & \xrightarrow{\rho_d} & C[L_0^{\odot},\dots,L_{i-1}^{\odot},L_i^{\bullet},L_{i+1},\dots,L_n] \\
\sigma_e \Downarrow & & \sigma_e \Downarrow \\
 & & \sigma_{e-1} \Downarrow \\
C[L_0^{\odot},\dots,L_{i-1}^{\odot},L_i^{\circ},L_{i+1},\dots,L_n] & \xrightarrow{\sigma_{e-1}} \xrightarrow{\rho_d} & C[L_0^{\odot},\dots,L_i^{\odot},L_{i+1},\dots,L_n]
\end{array}
$$

Every hole on the diagonal is filled with at most one ρ_d-redex (at the root of the context hole fillings) but because of $d < e$ with possibly many σ_e-redexes (properly inside of the fillings). There may or may not be an overlap between the ρ_d-step and a σ_e-step, but there can be at most one, which is due to the rules of $\boldsymbol{R}_{\nabla_\bullet}$.

Therefore σ_e contracts either an overlap and a number of nested redexes, or only nested redexes without an overlap. These constellations are depicted on the figure below. There is one ρ_d-redex and three σ_e-redexes. On the left, one of the σ_e-redexes overlaps with the ρ_d-redex while on the right all σ_e-redexes are strictly nested inside the ρ_d-redex.

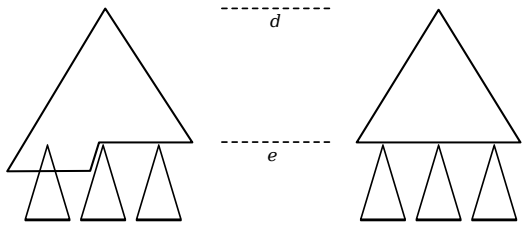

For the critical pairs due to a non-root overlap, and for all situations with nested redexes, we construct diagrams of the following shape, respectively:

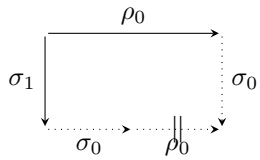

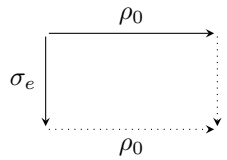

$e' \in \{e, e-1\}$

When lifted into a context of letrec-depth d both of the diagrams comply to the shape necessary for the diagonal tiles, but we need to be able to handle situations as on as on the left of the above figure, where both nested redexes as well as the overlapping redex are contracted. Firstly, since all σ-redexes occur at the same letrec-depth, it must hold that $d = 0$ and $e = 1$, which is due to the rules of $\boldsymbol{R}_{\triangledown\bullet}$. Secondly, none of the involved redex contractions affect any of the nested redexes except for duplicating or erasing them, which means that the residuals of the σ-steps after these steps are part of a parallel $\sigma_{e'}$-step (mind that we assume the reflexive closure of all steps). Or as a diagram:

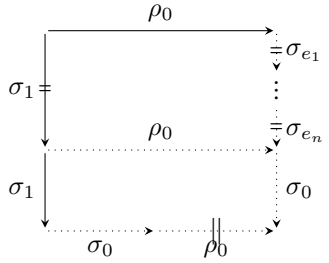

$e_i \in \{0, 1\}$

The diagram is composed from the previous two diagrams. A parallel version of the right one constitutes the top part, while the bottom part is an exact replica of the left one. The top part settles the portion arising from the nested

redexes, the bottom part settles the portion arising from the overlapping redex.

At last in order to fit that diagram into the scheme of the diagonal tiles the steps on the right have to be reordered such that σ_{e_i}-steps with $e_i = 1$ preceed σ_{e_i}-steps with $e_i = 0$. The reordering is viable because every σ_{e_i}-step takes place in its own residual of the σ_1-step from the left.

We conclude the proof by a comprehensive analysis all critical pairs that arise from non-root overlaps in $\boldsymbol{R}_{\triangledown\bullet}$ as well as the diagrams for joining nested redexes.

Diagrams for joining critical pairs

ρ_0/σ_0 critical pairs only arise for $\rho = \mathsf{letrec}$.

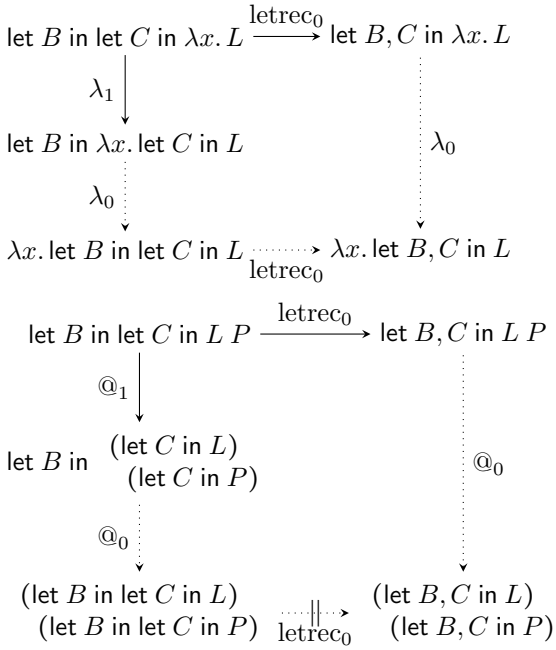

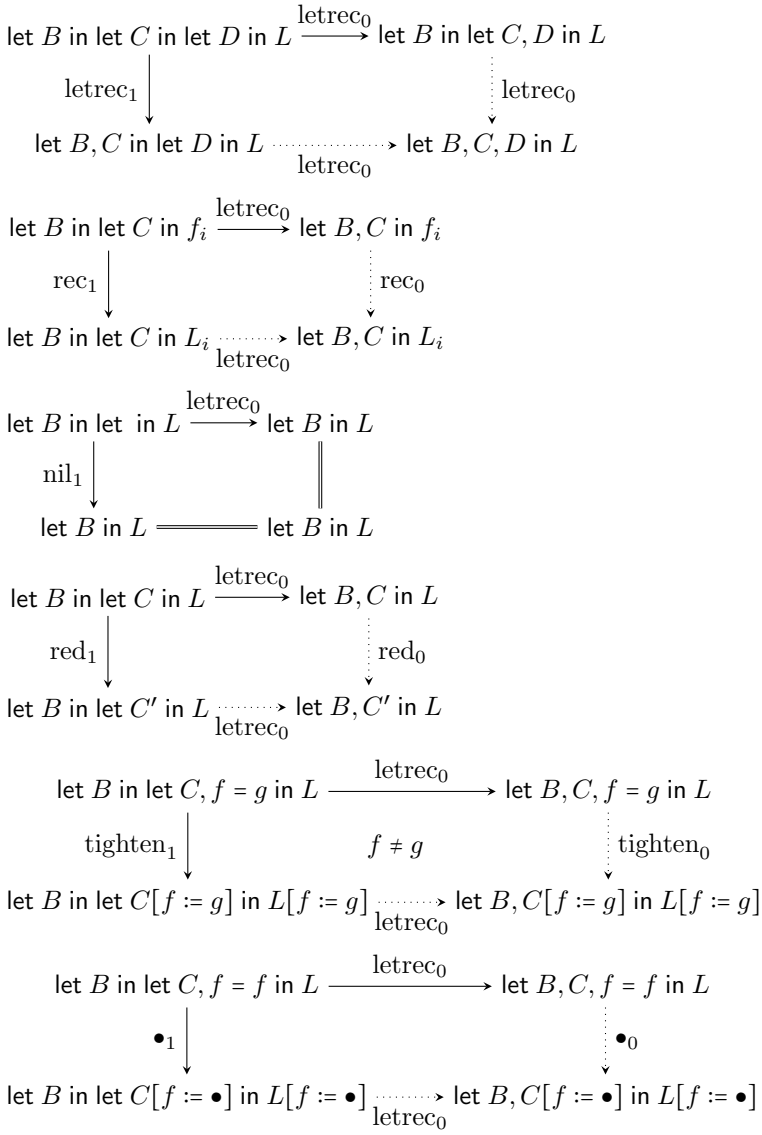
let B in let C in let D in L
letrec$_0$
let B in let C, D in L
letrec$_1$
letrec$_0$
let B, C in let D in L
letrec$_0$
let B, C, D in L
let B in let C in f_i
letrec$_0$
let B, C in f_i
rec$_1$
rec$_0$
let B in let C in L_i
letrec$_0$
let B, C in L_i
let B in let in L
letrec$_0$
let B in L
nil$_1$
let B in L
let B in L
let B in let C in L
letrec$_0$
let B, C in L
red$_1$
red$_0$
let B in let C' in L
letrec$_0$
let B, C' in L
let B in let $C, f = g$ in L
letrec$_0$
let $B, C, f = g$ in L
tighten$_1$
$f \neq g$
tighten$_0$
let B in let $C[f := g]$ in $L[f := g]$
letrec$_0$
let $B, C[f := g]$ in $L[f := g]$
let B in let $C, f = f$ in L
letrec$_0$
let $B, C, f = f$ in L
$\bullet_1$
$\bullet_0$
let B in let $C[f := \bullet]$ in $L[f := \bullet]$
letrec$_0$
let $B, C[f := \bullet]$ in $L[f := \bullet]$

Diagrams for joining nested redexes

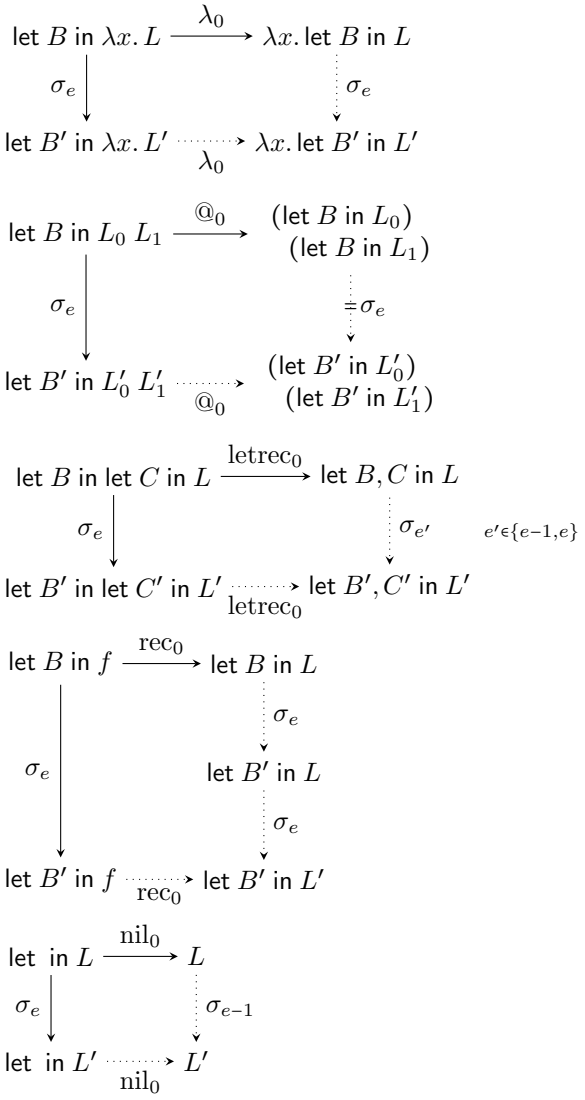

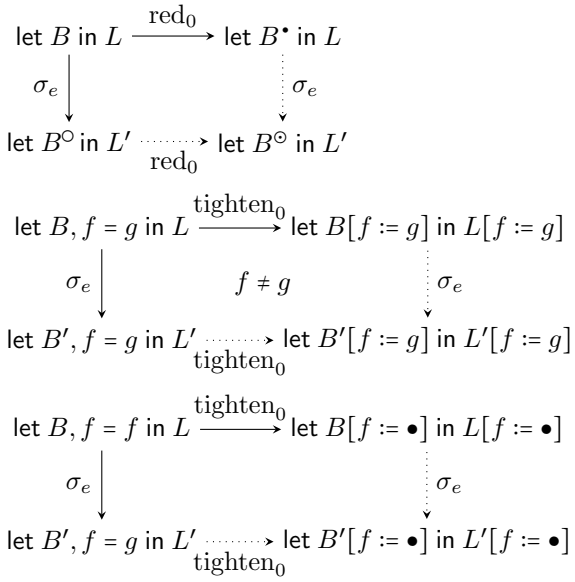
let B in L
$\xrightarrow{\mathrm{red}_0}$
let $B^\bullet$ in L
σ_e
σ_e
let $B^\bigcirc$ in L'
$\xrightarrow[\mathrm{red}_0]{}$
let $B^\odot$ in L'
let $B, f = g$ in L
$\xrightarrow{\mathrm{tighten}_0}$
let $B[f := g]$ in $L[f := g]$
σ_e
$f \neq g$
σ_e
let $B', f = g$ in L'
$\xrightarrow[\mathrm{tighten}_0]{}$
let $B'[f := g]$ in $L'[f := g]$
let $B, f = f$ in L
$\xrightarrow{\mathrm{tighten}_0}$
let $B[f := \bullet]$ in $L[f := \bullet]$
σ_e
σ_e
let $B', f = g$ in L'
$\xrightarrow[\mathrm{tighten}_0]{}$
let $B'[f := \bullet]$ in $L'[f := \bullet]$

□

Appendix D

Unfolding with a Single Rule

Here, we present an alternative rewriting system $\triangledown_{\mathbf{1}}$ for unfolding $\lambda_{\mathsf{letrec}}$-terms, which has only a single rule, similar to μ-unfolding. We will briefly argue that the system unfolds to the same λ-terms as $\boldsymbol{R}_{\triangledown}$. For simplicity we will not address the problem of meaningless bindings, which can be resolved with black holes and appropriate rewriting rules, as done for $\boldsymbol{R}_{\triangledown}$.

Notation D.25. Hereinafter, B stands for $f_1 = L_1, \ldots, f_n = L_n$. Furthermore $[f_i := L_i \mid 1 \le i \le n]$ stands for the substitution $[f_1 := L_1, \ldots, f_n := L_n]$.

Definition D.26 (One-rule rewriting system for unfolding $\lambda_{\mathsf{letrec}}$)**.** The rewriting system $\triangledown_{\mathbf{1}}$ consists of the single rule:

$$\mathsf{let}\ B\ \mathsf{in}\ L\ \to\ L[f_i := L_i[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \mid 1 \le i \le n]$$

In CRS notation:

$$\mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(L_1(\vec{f}),\ldots,L_n(\vec{f})\big),L_0(\vec{f})\big)\to L_0(L_1(\vec{B}),\ldots,L_n(\vec{B}))$$

$$\text{where } \vec{B} \text{ stands for } \left(\begin{array}{c}\mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(L_1(\vec{f}),\ldots,L_n(\vec{f}),L_1(\vec{f})\big)\big),\\ \ldots,\\ \mathsf{let}\big([\vec{f}]\,\mathsf{in}_n\big(L_1(\vec{f}),\ldots,L_n(\vec{f}),L_n(\vec{f})\big)\big)\end{array}\right)$$

§ D.27 (an even simpler approach?)**.** One might wonder why the nested substitutions are necessary, why a simpler rule as follows would not be adequate:

$$\mathsf{let}\ B\ \mathsf{in}\ L \to L[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n]$$

This rule is not sufficient because it cannot unfold a term $\mathsf{let}\ B\ \mathsf{in}\ L$ correctly if L is a variable bound in B; it reduces a term like $\mathsf{let}\ f = x\ \mathsf{in}\ f$ to itself.

Proposition D.28. $\triangledown_{\mathbf{1}}$ is confluent. (*Proof:* $\triangledown_{\mathbf{1}}$ is orthogonal.)

In order to argue that $\triangledown_{\mathbf{1}}$ reduces $\lambda_{\mathsf{letrec}}$-terms to the same infinite normal forms as $\boldsymbol{R}_{\triangledown}$ (let us write this as $\triangledown_{\mathbf{1}} \equiv \boldsymbol{R}_{\triangledown}$), we introduce an intermediate system $\triangledown_{\mathbf{2}}$ consisting of the rule from § D.27, and one additional rule from $\boldsymbol{R}_{\triangledown}$ to handle the problematic case.

Definition D.29 (Two-rule rewriting system for unfolding $\lambda_{\mathsf{letrec}}$)**.** The rewriting system $\triangledown_{\mathbf{2}}$ consists of the following two rules:

$$\begin{array}{l}\mathsf{let}\ B\ \mathsf{in}\ L \to L[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n] \\ \quad \text{(if } L \text{ is not a variable bound in } B\text{)} \\ \mathsf{let}\ B\ \mathsf{in}\ f_i \to \mathsf{let}\ B\ \mathsf{in}\ L_i\end{array}$$

Proposition D.30 ($\triangledown_{\mathbf{2}}$ is a refinement of $\triangledown_{\mathbf{1}}$)**.** $\to_{\triangledown_{\mathbf{1}}} \subseteq \twoheadrightarrow_{\triangledown_{\mathbf{2}}}$

Proof. Let us consider such a redex $\mathsf{let}\ B\ \mathsf{in}\ L$.

Case 1 (L is not a variable bound by B).

Case 1.1 (all the L_i below are not variables bound by B)

$$\begin{array}{lll}\mathsf{let}\ B\ \mathsf{in}\ L \to_{\triangledown_{\mathbf{2}}} & L[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n] \\ \twoheadrightarrow_{\triangledown_{\mathbf{2}}} & L[f_i := \mathsf{let}\ B\ \mathsf{in}\ L_i \mid 1 \le i \le n] \\ \twoheadrightarrow_{\triangledown_{\mathbf{2}}} & L[f_i := L_i[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \mid 1 \le i \le n] \\ \leftarrow_{\triangledown_{\mathbf{1}}} & \mathsf{let}\ B\ \mathsf{in}\ L\end{array}$$

Case 1.2 (all the L_i below are variables bound by B ($L_i = f_{j_i}$))

$$\begin{array}{lll}\mathsf{let}\ B\ \mathsf{in}\ L \to_{\triangledown_{\mathbf{2}}} & L[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n] \\ \twoheadrightarrow_{\triangledown_{\mathbf{2}}} & L[f_i := \mathsf{let}\ B\ \mathsf{in}\ L_i \mid 1 \le i \le n] \\ = & L[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_{j_i} \mid 1 \le i \le n] \\ = & L[f_i := f_{j_i}[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \mid 1 \le i \le n] \\ = & L[f_i := L_i[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \mid 1 \le i \le n] \\ \leftarrow_{\triangledown_{\mathbf{1}}} & \mathsf{let}\ B\ \mathsf{in}\ L\end{array}$$

Actually, the case distinction of **Case 1** into **Case 1.1** and **Case 2.2** is non-exhaustive. Instead, all mixtures of **Case 1.1** and **Case 2.2** have to be considered, where some of the L_i are bound by B and others are not. To write this out would however be merely a tedious exercise.

Case 2 (L is a variable bound by B ($L = f_i$)).

Case 2.1 (L_i below is not a variable bound by B)

$$\begin{aligned}
\mathsf{let}\ B\ \mathsf{in}\ f_i &\to_{\triangledown_{\mathbf{2}}} && \mathsf{let}\ B\ \mathsf{in}\ L_i \\
&\to_{\triangledown_{\mathbf{2}}} && L_i[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n] \\
&= && f_i[f_i := L_i[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \mid 1 \le i \le n] \\
&\leftarrow_{\triangledown_{\mathbf{1}}} && \mathsf{let}\ B\ \mathsf{in}\ f_i
\end{aligned}$$

Case 2.2 (L_i below is a variable bound by B ($L_i = f_j$))

$$\begin{aligned}
\mathsf{let}\ B\ \mathsf{in}\ f_i &\to_{\triangledown_{\mathbf{2}}} && \mathsf{let}\ B\ \mathsf{in}\ L_i \\
&= && \mathsf{let}\ B\ \mathsf{in}\ f_j \\
&= && f_j[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n] \\
&= && L_i[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n] \\
&= && f_i[f_i := L_i[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \mid 1 \le i \le n] \\
&\leftarrow_{\triangledown_{\mathbf{1}}} && \mathsf{let}\ B\ \mathsf{in}\ f_i
\end{aligned}$$

□

Proposition D.31 ($\triangledown_{\mathbf{2}}$ is a compatible refinement of $\triangledown_{\mathbf{1}}$)**.** $\triangledown_{\mathbf{2}}$ is a refinement of $\triangledown_{\mathbf{1}}$, and additionally it holds:

$$\forall M, N \in \mathrm{Ter}(\boldsymbol{\lambda}_{\mathsf{letrec}})\ (M \twoheadrightarrow_{\triangledown_{\mathbf{2}}} N \ \Rightarrow\ M \twoheadrightarrow_{\triangledown_{\mathbf{1}}} . \twoheadleftarrow_{\triangledown_{\mathbf{1}}} N)$$

Proof. Let us consider a $\to_{\triangledown_{\mathbf{2}}}$-redex $M = \mathsf{let}\ B\ \mathsf{in}\ L$ and the reduction $M \to_{\triangledown_{\mathbf{2}}} N$. We will show that $M \twoheadrightarrow_{\triangledown_{\mathbf{1}}} . \twoheadrightarrow_{\triangledown_{\mathbf{1}}} N$

Case 1 (L is not a variable bound by B).

$$M = \mathsf{let}\ B\ \mathsf{in}\ L \to_{\triangledown_{\mathbf{2}}} L[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n] = N$$

$$\begin{array}{ll} M \to_{\triangledown_{\mathbf{1}}} & L[f_i := L_i[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \mid 1 \le i \le n] \\ = & L[f_i := f_i[f_i := L_i[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \mid 1 \le i \le n] \mid 1 \le i \le n] \\ \leftarrow_{\triangledown_{\mathbf{1}}} & N \end{array}$$

Case 2 (L is a variable bound by B ($L = f_i$)).

$$M = \mathsf{let}\ B\ \mathsf{in}\ f_i \to_{\triangledown_{\mathbf{2}}} \mathsf{let}\ B\ \mathsf{in}\ L_i = N$$

$$\begin{array}{ll} M \to_{\triangledown_{\mathbf{1}}} & f_i[f_i := L_i[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \mid 1 \le i \le n] \\ = & L_i[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \\ \to_{\triangledown_{\mathbf{1}}} & L_i[f_j := f_j[f_i := L_i[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \mid 1 \le i \le n] \mid 1 \le j \le n] \\ = & L_i[f_i := L_i[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \mid 1 \le i \le n] \\ \leftarrow_{\triangledown_{\mathbf{1}}} & \mathsf{let}\ B\ \mathsf{in}\ L_i \end{array}$$

□

Proposition D.32. $\triangledown_{\mathbf{2}}$ is confluent.

Proof. This follows from proposition D.28 and proposition D.31 [44]. □

Proposition D.33 ($\triangledown_{\mathbf{1}} \equiv \triangledown_{\mathbf{2}}$)**.**

$$\forall L \in \mathrm{Ter}(\lambda_{\mathsf{letrec}})\ L \twoheadrightarrow^{!}_{\triangledown_{\mathbf{1}}} L' \iff L \twoheadrightarrow^{!}_{\triangledown_{\mathbf{2}}} L'$$

Proposition D.34 ($\boldsymbol{R}_{\triangledown} \equiv \triangledown_{\mathbf{2}}$)**.**

$$\forall L \in \mathrm{Ter}(\lambda_{\mathsf{letrec}})\ L \twoheadrightarrow^{!}_{\triangledown} L' \iff L \twoheadrightarrow^{!}_{\triangledown_{\mathbf{2}}} L'$$

Proof idea. Only a rudimentary proof idea is given here, i.e. that the property

$$\to_{\triangledown} \subseteq \twoheadrightarrow_{\triangledown_{\mathbf{2}}} . \twoheadleftarrow_{\triangledown_{\mathbf{2}}}$$

can be used to construct a coinductive proof. In order to convince ourselves that the property holds we only look at three cases here. In each of these cases the leftmost term reduces via $\to_{\triangledown}$ to the rightmost term.

Case: λ

$$\begin{array}{lll} \mathsf{let}\ B\ \mathsf{in}\ \lambda x.\, L & \to_{\triangledown_{\mathbf{2}}} & (\lambda x.\, L)[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n] \\ & = & \lambda x.\, L[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n] \quad \twoheadleftarrow_{\triangledown_{\mathbf{2}}} \lambda x.\, \mathsf{let}\ B\ \mathsf{in}\ L \end{array}$$

Case: @

$$\begin{array}{lll} \mathsf{let}\ B\ \mathsf{in}\ L\ P & \to_{\triangledown_{\mathbf{2}}} (L\ P)[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n] & \\ & = \begin{array}{l} (L[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n]) \\ (P[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n]) \end{array} & \twoheadleftarrow_{\triangledown_{\mathbf{2}}} \begin{array}{l} (\mathsf{let}\ B\ \mathsf{in}\ L) \\ (\mathsf{let}\ B\ \mathsf{in}\ P) \end{array} \end{array}$$

Case: letrec

$$\begin{array}{ll} & \mathsf{let}\ B\ \mathsf{in}\ \mathsf{let}\ C\ \mathsf{in}\ L \\ \to_{\triangledown_{\mathbf{2}}} & (\mathsf{let}\ C\ \mathsf{in}\ L)[f_i := \mathsf{let}\ B\ \mathsf{in}\ f_i \mid 1 \le i \le n] \\ \to_{\triangledown_{\mathbf{2}}} & (L[g_i := \mathsf{let}\ C\ \mathsf{in}\ g_i \mid 1 \le i \le n])[f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n] \\ = & L([g_i := \mathsf{let}\ C\ \mathsf{in}\ g_i \mid 1 \le i \le n] \cup [f_j := \mathsf{let}\ B\ \mathsf{in}\ f_j \mid 1 \le j \le n]) \\ \leftarrow_{\triangledown_{\mathbf{2}}} & \mathsf{let}\ B, C\ \mathsf{in}\ L \end{array}$$

□

Proposition D.35 ($\triangledown_{\mathbf{1}} \equiv \boldsymbol{R}_{\triangledown}$)**.**

$$\forall L \in \mathrm{Ter}(\lambda_{\mathsf{letrec}})\ L \twoheadrightarrow^{!}_{\triangledown_{\mathbf{1}}} L' \iff L \twoheadrightarrow^{!}_{\triangledown} L'$$

Proof. Follows from proposition D.34 and proposition D.33. □

Curriculum Vitae

birth: 20 July 1984 in Bretten, Germany

1990–1994 **Elementary school**, *Grundschule Nussbaum*

1994–2003 **Secondary school**, *Melanchthon-Gymnasium Bretten*
Major subjects: mathematics, physics

2002–2003 **Schülerstudium Informatik**, *Universität Karlsruhe*
As a 'student with exceptional learning achievements' I was given the opportunity to study computer science at the university during my last year of secondary school.

2003–2004 **Research assistance**, *Forschungszentrum Informatik, Karlsruhe*

2003–2009 **Study of computer science**, *Universität Karlsruhe*
Major subjects: system architecture, compiler construction and software engineering

2009 **Research semester**, *Chalmers University of Technology, Göteborg*
Diploma thesis under the supervision of the functional programming research group

2009-2013 **Doctoral research**, *Departement Informatica, Universiteit Utrecht*
Research and teaching assistance in the field of functional programming

Samenvatting in het Nederlands

(Summary in Dutch)

Dit proefschrift bevat resultaten uit onderzoek van de ongetypeerde λ-calculus (lees: ,,lambda calculus"). De λ-calculus is een formeel systeem dat in de jaren 1930 is ontworpen en is sindsdien onderzoeksonderwerp in informatica en filosofie:

- Filosofie > Formele Logica > Herschrijven > λ-Calculus
- Informatica > Programmeertalen en Compilerconstructie > Functionele Programmeertalen > λ-Calculus

Het onderzoek is voornamelijk gemotiveerd door de rol van de λ-calculus als basis van functionele programmeertalen. Het doel van het onderzoeksproject was om de uitvoering van in functionele programeertalen geschreven programma's efficiënter te maken door het verhogen van *sharing*. Het Engelse begrip sharing duidt het fenomeen aan dat een berekening *gedeeld* wordt, dat wil zeggen dat een waarde die op meerderen plekken nodig is niet iedere keer opnieuw wordt uitgerekend, maar dat die waarde slechts één keer wordt uitgerekend en dat het resultaat vervolgens op deze plekken wordt hergebruikt. Om dit te bereiken worden functionele talen doorgaans als graafherschrijfsysteem geïmplementeerd. In een graaf kan een knoop meerdere inkomende kanten hebben. Op die manier wordt die knoop en zijn opvolgers gedeeld. Maar vaak moet tijdens de uitvoering een gedeelde subgraaf worden ,,ontdeeld" of ,,ontvouwd". Hoe later en hoe voorzichtiger het ontdelen kan worden doorgevoerd des te hoger is de graad van sharing en de hierdoor bespaarde berekeningskosten. Er zijn verschillende evaluatiemodellen met verschillende graden van sharing. In dit proefschrift focussen wij niet op het dynamische delen tijdens de uitvoering maar op het statische delen *voor* de uitvoerig van het programma. Het delen tijdens de uitvoering is namelijk gebaseerd op de initiële graad van sharing van de graaf

die door de compiler van een programma wordt geconstrueerd. Deze graaf wordt vaak in een taal beschreven die gebaseerd is op de λ-*calculus met* letrec (of kort: $\lambda_{\mathsf{letrec}}$). Zo noemen we de λ-calculus waarvan de termen niet alleen normale λ-termen zijn maar ook termen met voorkomens van de zogenoemde letrec-constructie. De letrec-constructie maakt het mogelijk om een graafstructuur en dus sharing direct in de programmeercode uit te drukken.

In dit proefschrift bestuderen we sharing in $\lambda_{\mathsf{letrec}}$. Het centrale begrip voor hoe we een $\lambda_{\mathsf{letrec}}$-term beschouwen is de *ontvouwingssemantiek*: We beschouwen een $\lambda_{\mathsf{letrec}}$-term als een representatie van een (potientieel) oneindige term. De rechtvaardiging hiervoor is dat dit overeenkomt met hoe functionele talen geëvalueerd worden: voordat een β-reductie plaatsvindt wordt het deel van de graaf dat de redex bevat eerst ontvouwd en het ontvouwen van een $\lambda_{\mathsf{letrec}}$-term heeft (in de limiet) een (potentieel oneindige) λ-term als resultaat. We houden ons niet verder bezig met β-reductie maar onderzoeken uitsluitend ontvouwing van $\lambda_{\mathsf{letrec}}$-termen.

Hoofdstuk 0 introduceert het $\lambda_{\mathsf{letrec}}$-formalisme samen met een herschrijfsysteem om $\lambda_{\mathsf{letrec}}$-termen te ontvouwen. In de volgende drie hooftstukken bestuderen we dit herschrijfsysteem en beantwoorden o. a. de volgende vragen:

Welke oneindige λ-termen kunnen als eindige $\lambda_{\mathsf{letrec}}$-termen worden uitgedrukt? Met anderen woorden: hoe expressief is de taal $\boldsymbol{\lambda}_{\mathsf{letrec}}$? In hoofdstuk 1 karakteriseren we de verzameling van $\lambda_{\mathsf{letrec}}$-uitdrukbaren λ-termen, dus die λ-termen die de ontvouwing van eindige $\lambda_{\mathsf{letrec}}$-termen zijn.

Wat zijn goede graafrepresentaties voor $\lambda_{\mathsf{letrec}}$-uitdrukbare termen? De letrec-constructie dient ertoe om een graafstructuur uit te drukken. In hoofdstuk 2 vragen we ons af hoe precies deze grafen kunnen worden geformaliseerd en we identificeren een formalisatie die voor de volgende vragen nuttig blijkt.

Hoe kunnen we bepalen of twee $\lambda_{\mathsf{letrec}}$-termen dezelfde ontvouwing hebben? Hoe kunnen we van een $\lambda_{\mathsf{letrec}}$-term een equivalente maar zo compact mogelijke variante berekenen? Bestaat er zoiets als een zo compact mogelijke variant? In hoofdstuk 3 tonen we aan dat er voor iedere klasse van equivalente $\lambda_{\mathsf{letrec}}$-termen een maximaal compacte vorm bestaat. We ontwikkelen praktische en efficiënte methoden om deze vorm te berekenen en om te bepalen of twee $\lambda_{\mathsf{letrec}}$-termen dezelfde ontvouwing hebben.

Lay Summary

This summary is intended to give the casual reader an idea what this thesis is about.

In this thesis I study a very specific subject on the intersection of the fields of computer science and philosophy, namely a formal system called the λ-calculus (read 'lambda calculus'). The λ-calculus can be placed into these two fields of science as follows:

- Philosophy > Formal Logic > Rewriting Systems > λ-Calculus
- Computer Science > Programming Languages and Compiler Construction > Functional Programming Languages > λ-Calculus

In order to convey the meaning of the thesis title, let us first establish what computers and programming languages are all about.

A computer is a machine for processing information. It has input channels (like a touchpad or a microphone) to receive data; it processes the data and performs computations; and it has output channels (like a display or a connector to another device) to convey the result of the computation, by making the result visible, or carrying out some mechanical activity.

The behaviour of computers is determined by their programming. They are *programmable*, meaning their behaviour can be changed by a programmer. But moreover, a computer *requires* programming to function; there is no intelligence inherent to the device itself; it is the programmer that imbues the machine with his ideas through programming.

The programmer's ideas (say, a method to increase the contrast in a photograph) have to be expressed such that it can be understood by the computer. The language that a computer understands is its *machine language* which is usually a very simple language and which is different for each type of

computer. It consists of a collection of basic instructions that, when carried out by the computer, each make very small changes to the computer's memory. It is difficult and cumbersome for a programmer to express sophisticated ideas using a machine language.

Today, programmers do not need to program directly with machine languages, but have a vast number of *programming languages* at their disposal. Programming languages are languages designed to make the task of programming more efficient and convenient, by giving programmers useful, more powerful means to encode their ideas. A program called *compiler* then translates the code written in a programming language into machine language, which subsequently can be executed by the machine. A very simple compiler that uses the most straightforward way to translate a programming language into machine language produces very inefficient machine language, which leads to a slower execution of the program. This is especially the case for more complex programming languages, which provide the programmer with a higher level of convenience. Modern compilers produce more efficent machine language by analysing the original program and then translating it in a more sophisticated way. Such a compiler is called an *optimising* compiler.

This thesis focusses on one specific sub-class of programming languages called *functional programming languages*. They differ from *imperative programming languages*, which form the mainstream of today's programming languages. Code written in an imperative programming language can be regarded as a list of instructions which are executed sequentially, one by one. Code written in a functional programming languages, on the other hand, more resembles a set of mathematical equations, by which the result of the computation is defined.

While functional programming has for decades led a niche existence mostly limited to the academic world, its merits and its elegance are being ever more recognised, and is steadily gaining ground. Today functional programming languages are employed in commerce and in industry, and modern imperative programming languages incorporate more and more features from functional programming languages.

Functional programming has its roots in the λ-calculus, a formal system developed in the 1930s. It is a *rewriting system*; that means it acts on a specific kind of formal expressions (λ-terms) which can be rewritten in a rewriting step to a different term by a fixed set of rewriting rules. It is a sort of minimalistic programming language; the algorithm and the input are expressed as a λ-term; the computation consists of repeatedly rewriting the term according to the rewriting rules of the λ-calculus until it cannot be rewritten any further; the term one obtains in the end is the result of the computation.

There has been a multitude of extensions of the λ-calculus. One such extension is the inclusion of a language construct called letrec, which we call the 'λ-calculus with letrec' or in short $\lambda_{\mathsf{letrec}}$. The letrec is a syntactic element which allows for function definitions. Function definitions are equations which bind terms to names. A bound term can be referenced by its name and thus used multiple times at different places, thereby introducing 'sharing'. In the course of performing a computation in $\lambda_{\mathsf{letrec}}$ the definitions need to be *unfolded*. Unfolding is the process of replacing occurrences of function names by their definition. As function names may occur within their own definition unfolding may go on forever resulting an infinite term. This thesis is all about the letrec-construct and unfolding; it tackles questions like:

- Which are the infinite terms that can be obtained by unfolding of finite $\lambda_{\mathsf{letrec}}$-terms?
- When do two given $\lambda_{\mathsf{letrec}}$-terms have the same unfolding?
- How can we find a maximally compact form of a given $\lambda_{\mathsf{letrec}}$-term?
- What is a suitable graph representation for $\lambda_{\mathsf{letrec}}$-terms?

These results provide compilers with further opportunities for analysis and for a more optimised translation which might speed up the execution of programs written in a functional programming languages. Furthermore they may also help reasoning about programs and thus promote further research about functional programming languages.

Acknowledgements

First and foremost my thanks go to Clemens Grabmayer: Danke für die Zusammenarbeit und viele interessante Diskussionen; das war eine sehr lehrreiche Zeit für mich und hat meinen Horizont sehr erweitert. But of course also many thanks to my promotors Doaitse Swierstra and Vincent van Oostrom: naast heel stimulerende gesprekken en uw vakkunde kon ik vooral het vertrouwen waarderen dat in mij is gesteld om zelfstandig mijn onderzoeksdoelen te kiezen en na te gaan. I am very grateful to my research group, the Centre for Software Technology of Utrecht University for offering me a workplace beyond the duration of my contract. Furthermore I thank my parents for their substantial support and my sister, Nora Rochel for the title page.

I want to thank the reading committee and my opponents for their comments and for a challenging and pleasant defense.